arXiv:1402.3314v2 [cs.LO] 16 Jul 2014

# Distributed synthesis for acyclic architectures

Anca Muscholl
Université de Bordeaux and Igor Walukiewicz
CNRS, Université de Bordeaux

July 21, 2014

**Abstract**

The distributed synthesis problem is about constructing correct distributed systems, i.e., systems that satisfy a given specification. We consider a slightly more general problem of distributed control, where the goal is to restrict the behavior of a given distributed system in order to satisfy the specification. Our systems are finite state machines that communicate via rendez-vous (Zielonka automata). We show decidability of the synthesis problem for all $\omega$-regular local specifications, under the restriction that the communication graph of the system is acyclic. This result extends a previous decidability result for a restricted form of local reachability specifications.

## 1 Introduction

Synthesizing distributed systems from specifications is an attractive objective, since distributed systems are notoriously difficult to get right. Unfortunately, there are very few known decidable frameworks for distributed synthesis. We study a framework for synthesis of open systems that is based on rendez-vous communication and causal memory. In particular, causal memory implies that specifications can talk about when a communication takes place, but cannot limit information that is transmitted during communication. This choice is both realistic and avoids some pathological reasons for undecidability. We show a decidability result for acyclic communication graphs and local $\omega$-regular specifications.

Instead of synthesis we actually work in the more general framework of distributed control. Our setting is a direct adaptation of the supervisory control framework of Ramadge and Wonham [15]. In this framework we are given a plant (a finite automaton) where some of the actions are uncontrollable, and a specification, and the goal is to construct a controller (another finite automaton) such that its product with the plant satisfies the specification. The controller is not allowed to block uncontrollable actions, in other words, in every state there is a transition on each uncontrollable action. The controlled plant has less behaviors, resulting from restricting controllable actions of the plant. In our case the formulation is exactly the same, but we consider Zielonka automata instead of finite automata, as plants and controllers. Considering parallel devices, as Zielonka automata, in the standard definition of control gives an elegant formulation of the distributed control problem.

Zielonka automata [17, 12] are by now a well-established model of distributed computation. Such a device is an asynchronous product of finite-state processes synchronizing on shared actions. Asynchronicity means that processes can progress at different speed. The synchronization on shared actions allows the synchronizing processes to exchange information, in particular the controllers can transfer control information with each synchronization. This model can encode some common synchronization primitives available on modern multi-core processors for implementing concurrent data structures, like compare-and-swap.

We show decidability of the control problem for Zielonka automata where the communication graph is acyclic: a process can communicate (synchronize) with its parent and its children. Our specifications are conjunctions of $\omega$-regular specifications for each of the component processes. We allow uncontrollable communication actions – the only restriction is that all communication actions must be binary. Uncontrollable communications give a big flexibility, for instance it is possible to model asymmetric situations where communication can be refused by one partner, but not by the other one.

Our result extends [5] that showed decidability for a restricted form of local reachability objectives (blocking final states). We still get the same complexity as in [5]: non-elementary in general, and EXPTIME for architectures of depth 1. Covering all $\omega$-regular objectives allows to express fairness constraints but at the same time introduces important technical obstacles. Indeed, for our construction to work it is essential that we enrich the framework by uncontrollable synchronization actions. This makes a separation into controllable and uncontrollable states impossible. In consequence, we are lead to abandon the game metaphor, to invent new arguments, and to design a new proof structure.

Most research on distributed synthesis and control has been done in the setting proposed by Pnueli and Rosner [14]. This setting is also based on shared-variable communication, however it does not allow to pass additional information between processes. So their model leads to partial information games, and decidability of synthesis holds only for very restricted architectures [8, 9, 2]. While specifications leading to undecidability are very artificial, no elegant solution to eliminate them exists at present. The synthesis setting is investigated in [9] for local specifications, meaning that each process has its own, linear-time specification. For such specifications, it is shown that an architecture has a decidable synthesis problem if and only if it is a sub-architecture of a pipeline with inputs at both endpoints. More relaxed variants of synthesis have been proposed, where the specification does not fully describe the communication of the synthesized system. One approach consists in adding communication in order to combine local knowledge, as proposed for example in [6]. Another approach is to use specifications only for describing external communication, as done in [4] on strongly connected architectures where processes communicate via signals.

Apart from [5], two closely related decidability results for synthesis with causal memory are known, both of different flavor than ours. The first one [3] restricts the alphabet of actions: control with reachability condition is decidable for co-graph alphabets. This restriction excludes among others client-server architectures, which are captured by our setting. The second result [10] shows decidability by restricting the plant: roughly speaking, the restriction says that every process can have only bounded missing knowledge about the other processes, unless they diverge (see also [13] that shows a doubly exponential upper

bound). The proof of [10] goes beyond the controller synthesis problem, by coding it into monadic second-order theory of event structures and showing that this theory is decidable when the criterion on the plant holds. Unfortunately, very simple plants have a decidable control problem but undecidable MSO-theory of the associated event structure. Game semantics and asynchronous games played on event structures are considered in [11]. More recent work [7] considers games on event structures and shows a Borel determinacy result for such games under certain restrictions.

*Overview.* In Section 2 we state our control problem, and in Section 3 we give the main lines of the proof, that works by a reduction of the number of processes. In Section 3.2 we show that we may assume for the process that is eliminated that there is a bound on the number of local actions it can perform between consecutive synchronizations with its parent. In Section 3.3 we present the reduction, and in Sections 3.4, 3.5 we show the correctness of the construction.

# 2 Control for Zielonka automata

In this section we introduce our control problem for Zielonka automata, adapting the definition of supervisory control [15] to our model.

A Zielonka automaton [17, 12] is a simple distributed finite-state devices. Such an automaton is a parallel composition of several finite automata, called *processes*, synchronizing on shared actions. There is no global clock, so between two synchronizations, two processes can do a different number of actions. Because of this, Zielonka automata are also called asynchronous automata.

A *distributed action alphabet* on a finite set $\mathbb{P}$ of processes is a pair $(\Sigma, dom)$, where $\Sigma$ is a finite set of *actions* and $dom : \Sigma \to (2^{\mathbb{P}} \setminus \emptyset)$ is a *location function*. The location $dom(a)$ of action $a \in \Sigma$ comprises all processes that need to synchronize in order to perform this action. Actions from $\Sigma_p = \{a \in \Sigma \mid p \in dom(a)\}$ are called *p-actions*. We write $\Sigma_p^{loc} = \{a \mid dom(a) = \{p\}\}$ for the set of *local* actions of $p$.

A (deterministic) *Zielonka automaton* $\mathcal{A} = \langle \{S_p\}_{p\in\mathbb{P}}, s_{in}, \{\delta_a\}_{a\in\Sigma}\rangle$ is given by:

- for every process $p$ a finite set $S_p$ of (local) states,
- the initial state $s_{in} \in \prod_{p\in\mathbb{P}} S_p$,
- for every action $a \in \Sigma$ a partial transition function $\delta_a : \prod_{p\in dom(a)} S_p \stackrel{\cdot}{\to} \prod_{p\in dom(a)} S_p$ on tuples of states of processes in $dom(a)$.

**Example 1** Boolean multi-threaded programs with shared variables can be modeled as Zielonka automata. As an example we describe the translation for the *compare-and-swap* (CAS) instruction. This instruction has 3 parameters: $\mathsf{CAS}(x$: variable; *old*, *new*: int). Its effect is to return the value of $x$ and at the same time set the value of $x$ to *new*, but only if the previous value of $x$ was equal to *old*. The compare-and-swap operation is a widely used primitive in implementations of concurrent data structures, and has hardware support in most contemporary multiprocessor architectures.

Suppose that we have a thread $t$, and a shared variable $x$ that is accessed by a CAS operation in $t$ via $y := \mathsf{CAS}_x(i, k)$. So $y$ is a local variable of $t$. In

the Zielonka automaton we will have one process modeling thread $t$ and one process for variable $x$. The states of $t$ will be valuations of local variables. The states of $x$ will be the values $x$ can take. The CAS instruction above becomes a synchronization action. We have the following two types of transitions on this action:

Notice that in state $s'$, we have $y = i$, whereas in $s''$, we have $y = j$.

For convenience, we abbreviate a tuple $(s_p)_{p\in P}$ of local states by $s_P$, where $P \subseteq \mathbb{P}$. We also talk about $S_p$ as the set of *p-states*.

A Zielonka automaton can be seen as a sequential automaton with the state set $S = \prod_{p\in\mathbb{P}} S_p$ and transitions $s \xrightarrow{a} s'$ if $(s_{dom(a)}, s'_{dom(a)}) \in \delta_a$, and $s_{\mathbb{P}\setminus dom(a)} = s'_{\mathbb{P}\setminus dom(a)}$. So the states of this automaton are the tuples of states of the processes of the Zielonka automaton. For a process $p$ we will talk about the *p-component* of the state. A run of $\mathcal{A}$ is a finite or infinite sequence of transitions starting in $s_{in}$. Since the automaton is deterministic, a run is determined by the sequence of labels of the transitions. We will write $run(u)$ for the run determined by the sequence $u \in \Sigma^\infty$. Observe that $run(u)$ may be undefined since the transition function of $\mathcal{A}$ is partial. We will also talk about the projection of the run on component $p$, denoted $run_p(u)$, that is the projection on component $p$ of the subsequence of the run containing the transitions involving $p$. We will assume that every local state of $\mathcal{A}$ occurs in some run. For finite $w$ let $state(w)$ be the last state in $run(w)$. By $dom(u)$ we denote the union of $dom(a)$, for all $a \in \Sigma$ occurring in $u$.

We will be interested in maximal runs of Zielonka automata. For parallel devices the notion of a maximal run is not that evident, as one may want to impose some fairness conditions. We settle here for a minimal sensible fairness requirement. It says that a run is maximal if processes that have only finitely many actions in the run cannot perform any additional action.

**Definition 2 (Maximal run)** For a word $w \in \Sigma^\infty$ such that $run(w)$ is defined, we say that $run(w)$ is *maximal* if there is no decomposition $w = uv$, and no action $a \in \Sigma$ such that $dom(v) \cap dom(a) = \emptyset$ and $run(uav)$ is defined.

Automata can be equipped with a *correctness condition*. We prefer to talk about correctness condition rather than acceptance condition since we will be interested in the set of runs of an automaton rather than in the set of words it accepts. We will consider local regular correctness conditions: every process has its own correctness condition $Corr_p$. A run of $\mathcal{A}$ is *correct* if for every process $p$, the projection of the run on the transitions of $\mathcal{A}_p$ is in $Corr_p$. Condition $Corr_p$ is specified by a set $T_p \subseteq S_p$ of terminal states and an $\omega$-regular set $\Omega_p \subseteq (S_p \times \Sigma_p \times S_p)^\omega$. A sequence $(s_p^0, a_0, s_p^1)(s_p^1, a_1, s_p^2)\ldots$ satisfies $Corr_p$ if either: (i) it is finite and ends with a state from $T_p$, or (ii) it is infinite and belongs to $\Omega_p$. At this stage the set of terminal states $T_p$ may look unnecessary, but it will simplify our constructions later.

Finally, we will need the notion of *synchronized product* $\mathcal{A}\times\mathcal{C}$ of two Zielonka automata. For $\mathcal{A} = \langle\{S_p\}_{p\in\mathbb{P}}, s_{in}, \{\delta_a^A\}_{a\in\Sigma}\rangle$ and $\mathcal{C} = \langle\{C_p\}_{p\in\mathbb{P}}, c_{in}, \{\delta_a^C\}_{a\in\Sigma}\rangle$

let $\mathcal{A} \times \mathcal{C} = \langle \{S_p \times C_p\}_{p \in \mathbb{P}}, (s_{in}, c_{in}), \{\delta_a^\times\}_{a \in \Sigma} \} \rangle$ where there is a transition from $(s_{dom(a)}, c_{dom(a)})$ to $(s'_{dom(a)}, c'_{dom(a)})$ in $\delta_a^\times$ iff $(s_{dom(a)}, s'_{dom(a)}) \in \delta_a^A$ and $(c_{dom(a)}, c'_{dom(a)}) \in \delta_a^C$.

To define the control problem for Zielonka automata we fix a distributed alphabet $\langle \mathbb{P}, dom : \Sigma \to (2^{\mathbb{P}} \setminus \emptyset) \rangle$. We partition $\Sigma$ into the set of *system actions* $\Sigma^{sys}$ and *environment actions* $\Sigma^{env}$. Below we will introduce the notion of controller, and require that it does not block environment actions. For this reason we speak about *controllable/uncontrollable* actions when referring to system/environment actions. We impose three simplifying assumptions: (1) All actions are at most binary ($|dom(a)| \leq 2$ for every $a \in \Sigma$); (2) every process has some controllable action; (3) all controllable actions are local. Among the three conditions only the first one is indeed a restriction of our setting. The other two are not true limitations, in particular controllable shared actions can be simulated by a local controllable choice, followed by non-controllable local or shared actions (see Proposition 6).

**Definition 3 (Controller, Correct Controller)** A *controller* is a Zielonka automaton that cannot block environment (uncontrollable) actions. In other words, from every state every environment action is possible: for every $b \in \Sigma^{env}$, $\delta_b$ is a total function. We say that a controller $\mathcal{C}$ *is correct for* a plant $\mathcal{A}$ if all maximal runs of $\mathcal{A} \times \mathcal{C}$ satisfy the correctness condition of $\mathcal{A}$.

Recall that an action is possible in $\mathcal{A} \times \mathcal{C}$ iff it is possible in both $\mathcal{A}$ and $\mathcal{C}$. By the above definition, environment actions are always possible in $\mathcal{C}$. The major difference between the controlled system $\mathcal{A} \times \mathcal{C}$ and and $\mathcal{A}$ is that the states of $\mathcal{A} \times \mathcal{C}$ carry the additional information computed by $\mathcal{C}$, and that $\mathcal{A} \times \mathcal{C}$ may have less behaviors, resulting from disallowing controllable actions by $\mathcal{C}$.

The correctness of $\mathcal{C}$ means that all the runs of $\mathcal{A}$ that are *allowed* by $\mathcal{C}$ are correct. In particular, $\mathcal{C}$ does not have a correctness condition by itself. Considering only maximal runs of $\mathcal{A} \times \mathcal{C}$ imposes some minimal fairness conditions: for example it implies that if a process can do a local action almost always, then it will eventually do some action.

**Definition 4 (Control problem)** Given a distributed alphabet $\langle \mathbb{P}, dom : \Sigma \to (2^{\mathbb{P}} \setminus \emptyset) \rangle$ together with a partition of actions $(\Sigma^{sys}, \Sigma^{env})$, and given a Zielonka automaton $\mathcal{A}$ over this alphabet, find a controller $\mathcal{C}$ over the same alphabet such that $\mathcal{C}$ is correct for $\mathcal{A}$.

The important point in our definition is that the controller has the same distributed alphabet as the automaton it controls, in other words the controller is not allowed to introduce additional synchronizations between processes.

**Example 5** We give an example showing how causal memory works and helps to construct controllers. Consider an automaton $\mathcal{A}$ with 3 processes: $p$, $q$, $r$. We would like to control it so that the only two possible runs of $\mathcal{A}$ are the following:

So $p$ and $q$ should synchronize on $\alpha$ when action $a$ happened before $b$, otherwise $q$ and $r$ should synchronize on $\beta$. Communication actions are uncontrollable, but the transitions of $\mathcal{A}$ are such that there are local controllable actions $c$ and $d$ that enable communication on $\alpha$ and $\beta$ respectively. So the controller should block either $c$ or $d$ depending on the order between $a$ and $b$. The transitions of $\mathcal{A}$ are as follows

$$\begin{array}{cccc} \delta_a(p_0,q_0)=(p_1,q_1) & \delta_a(p_0,q_1)=(p_1,q_2) & \delta_b(q_0,r_0)=(q_1,r_1) & \delta_b(q_1,r_0)=(q_2,r_1) \\ \delta_c(p_1)=p_2 & \delta_\alpha(p_2,q_2)=(p_3,q_3) & \delta_d(r_1)=r_2 & \delta_\beta(q_2,r_2)=(q_3,r_3) \end{array}$$

These transitions allow the two behaviors depicted above but also two unwanted ones, as say, when $a$ happens before $b$ and then we see $\beta$. Clearly, the specification of the desired behaviors can be formulated as a local condition on $q$. So by encoding some information in states of $q$ this condition can be expressed by a set of terminal states $T_q$. We will not do this for readability.

The controller $\mathcal{C}$ for $\mathcal{A}$ will mimic the structure of $\mathcal{A}$: for every state of $\mathcal{A}$ there will be in $\mathcal{C}$ a state with over-line. So, for example, the states of $q$ in $\mathcal{C}$ will be $\overline{q}_0,\ldots,\overline{q}_3$. Moreover $\mathcal{C}$ will have two new states $\underline{p_1}$ and $\underline{r_1}$. The transitions will be

$$\begin{array}{cccc} \delta_a(\overline{p}_0,\overline{q}_0)=(\overline{p}_1,\overline{q}_1) & \delta_a(\overline{p}_0,\overline{q}_1)=(\underline{p_1},\overline{q}_2) & \delta_b(\overline{q}_0,\overline{r}_0)=(\overline{q}_1,\overline{r}_1) & \delta_b(\overline{q}_1,\overline{r}_0)=(\overline{q}_2,\underline{r_1}) \\ \delta_c(\overline{p}_1)=\overline{p}_2 & \delta_c(\underline{p_1})=\bot & \delta_d(\overline{r}_1)=\overline{r}_2 & \delta_d(\underline{r_1})=\bot \\ \delta_\alpha(\overline{p}_2,\overline{q}_2)=(\overline{p}_3,\overline{q}_3) & & \delta_\beta(\overline{q}_2,\overline{r}_2)=(\overline{q}_3,\overline{r}_3) & \end{array}$$

Observe that $c$ is blocked in $\underline{p_1}$, and so is $d$ from $\underline{r_1}$. It is easy to verify that the runs of $\mathcal{A}\times\mathcal{C}$ are as required, so $\mathcal{C}$ is a correct controller for $\mathcal{A}$. (Actually the definition of a controller forces us to make transitions of $\mathcal{C}$ total on uncontrollable actions. We can do it in arbitrary way as this will not add new behaviors to $\mathcal{A}\times\mathcal{C}$.)

This example shows several phenomena. The states of $\mathcal{C}$ are the states of $\mathcal{A}$ coupled with some additional information. We formalize this later under a notion of covering controller. We could also see above a case where a communication is decided by one of the parties. Processes $p$, thanks to a local action, can decide if it wants to communicate via $\alpha$, but process $q$ has to accept $\alpha$ always. This shows the flexibility given by uncontrollable communication actions. Finally, we could see information passing during communication. In $\mathcal{C}$ process $q$ passes to $p$ and $r$ information about its local state (transitions on $a$ and on $b$).

We end the section by showing the assumption that controllable actions are local, is not a restriction.

**Proposition 6** The control problem for Zielonka automata where communication actions may be controllable, reduces to the setting where controllable actions are all local.

**Proof**
We start with an automaton $\mathcal{A}$ over a distributed alphabet $\langle\Sigma, dom\rangle$ and a correct covering controller $\mathcal{C}$. We define first a new automaton $\mathcal{A}'$ over an extended distributed alphabet $\langle\Sigma', dom'\rangle$ with $\Sigma'=\Sigma\cup\{ch(A)\mid A\subseteq\Sigma_p^{sys} \text{ for some } p\in\mathbb{P}\}$. All new actions are local: $dom'(ch(A))=\{p\}$ if $A\subseteq\Sigma_p^{sys}$; the domain

of other actions do not change. What changes is that all old actions become uncontrollable, and the only controllable actions in $\Sigma'$ are those of the form $ch(A)$.

- The set of $p$-states of $\mathcal{A}'$ is the set of $p$-states of $\mathcal{A}$, plus some new states of the form $\langle s_p, A\rangle$ where $s_p$ is a $p$-state of $\mathcal{A}$ and $A \subseteq \Sigma_p^{sys}$.
- For every old $p$-state $s_p$ we delete all outgoing controllable transitions and add
$$s_p \xrightarrow{ch(A)} \langle s_p, A\rangle\,,$$
for every set $A$ of controllable actions enabled in $s_p$. From $\langle s_p, A\rangle$ we put in $\mathcal{A}'$ transitions as follows. If $a \in A$ is local then we have $\langle s_p, A\rangle \xrightarrow{a} s'_p$ whenever $s_p \xrightarrow{a} s'_p$ in $\mathcal{A}$. If $a \in A \cap B$ and $dom(a) = \{p, p'\}$ then we have $(\langle s_p, A\rangle, \langle s_{p'}, B\rangle) \xrightarrow{a} (s'_p, s'_{p'})$ whenever $(s_p, s_{p'}) \xrightarrow{a} (s'_p, s'_{p'})$ in $\mathcal{A}$.
- The correctness condition of $\mathcal{A}'$ is a straightforward modification of the one of $\mathcal{A}$.

Assume first that $\mathcal{C}'$ is a correct covering controller for $\mathcal{A}'$. From $\mathcal{C}'$ we define the automaton $\mathcal{C}$ over the same sets of states, by modifying slightly the transitions as follows. Suppose that $c \xrightarrow{ch(A)} d$ is a (local) transition in $\mathcal{C}'_p$. Since $\mathcal{C}'$ is covering we have a transition of the form $s_p = \pi'(c) \xrightarrow{ch(A)} \pi'(d) = \langle s_p, A\rangle$ in $\mathcal{A}'$. Let $a \in A$ be local. Since $a$ is uncontrollable in $\mathcal{A}'$ and $\langle s_p, A\rangle \xrightarrow{a} s'_p$ (for some $s'_p$) we must also have $d \xrightarrow{a} e$ for some state $e$ of $\mathcal{C}'_p$, since $\mathcal{C}'$ is covering. We delete $c \xrightarrow{ch(A)} d$ from $\mathcal{C}'$ and replace $d \xrightarrow{a} e$ by $c \xrightarrow{a} e$. If $a$ is shared by $p, p'$, let us consider some transition $c' \xrightarrow{ch(B)} d'$ with $a \in B$ in $\mathcal{C}'_{p'}$. Since $a$ is uncontrollable in $\mathcal{A}'$ we find again some transition $(d, d') \xrightarrow{a} (e, e')$ in $\mathcal{C}'$. We replace then $(d, d') \xrightarrow{a} (e, e')$ by $(c, c') \xrightarrow{a} (e, e')$ in $\mathcal{C}$. Of course, this is done in parallel for all transitions labeled by some $ch(A)$. It is immediate that $\mathcal{C}$ is covering $\mathcal{A}$, by taking $\pi = \pi'$. Maximal runs of $\mathcal{C}$ map to maximal runs of $\mathcal{C}'$ and thus satisfy the correctness condition for $\mathcal{A}$.

Conversely, given a correct covering controller $\mathcal{C}$ for $\mathcal{A}$ we define $\mathcal{C}'$ for $\mathcal{A}'$. Local $p$-states of $\mathcal{C}'$ are those of $\mathcal{C}$, plus additional states of the form $c_A$, where $c$ is a $p$-state of $\mathcal{C}$ and $A \subseteq \Sigma_p^{sys}$. Consider any $p$-state $c$ of $\mathcal{C}$, and let $A$ be the set of controllable actions enabled in $c$ (a communication action $a$ with $dom(a) = \{p, p'\}$ is enabled in $c$ if there exists some $p'$-state $c'$ and an $a$-transition from $(c, c')$). We replace all controllable transitions from $c$ by one (local) controllable transition $c \xrightarrow{ch(A)} c_A$, plus some uncontrollable transitions. If $a \in A$ is local, then we add the uncontrollable transitions $c_A \xrightarrow{a} d$ whenever $c \xrightarrow{a} d$ in $\mathcal{C}$. If $dom(a) = \{p, p'\}$, $(c, c') \xrightarrow{a} (d, d')$ in $\mathcal{C}$, and $B$ is the set of controllable actions enabled in the $p'$-state $c'$, then we replace $(c, c') \xrightarrow{a} (d, d')$ by $(c_A, c'_B) \xrightarrow{a} (d, d')$. Extending $\pi$ by $\pi'(c_A) = \langle \pi(c), A\rangle$ shows that $\mathcal{C}'$ is a covering controller for $\mathcal{A}'$. Maximal runs of $\mathcal{C}'$ satisfy the acceptance condition as for $\mathcal{C}$. □

# 3 Decidability for acyclic architectures

In this section we present the main result of the paper. We show the decidability of the control problem for Zielonka automata with acyclic architecture. A *communication architecture* of a distributed alphabet is a graph where nodes are processes and edges link processes that have common actions. An *acyclic architecture* is one whose communication graph is acyclic. For example, the communication graph of the alphabet from the example on page 5 is a tree with the root $q$ and two successors, $p$ and $r$.

**Theorem 7** *The control problem for Zielonka automata over distributed alphabets with acyclic architecture is decidable. If a controller exists, then it can be effectively constructed.*

The remaining of this section is devoted to the outline of the proof of Theorem 7. This proof works by induction on the number $|\mathbb{P}|$ of processes in the automaton. A Zielonka automaton over a single process is just a finite automaton, and the control problem is then just the standard control problem as considered by Ramadge and Wonham but extended to all $\omega$-regular conditions [1]. If there are several processes that do not communicate, then we can solve the problem for each process separately.

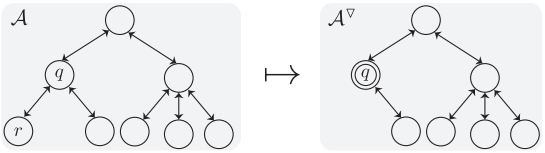


Figure 1: Eliminating process $r$: $r$ is glued with $q$.

Otherwise we choose a leaf process $r$ and its parent $q$, and construct a new plant $\mathcal{A}^\triangledown$ over $\mathbb{P}\setminus\{r\}$.We will show that the control problem for $\mathcal{A}$ has a solution iff the one for $\mathcal{A}^\triangledown$ does. Moreover, for every solution for $\mathcal{A}^\triangledown$ we will be able to construct a solution for $\mathcal{A}$.

For the rest of this section let us fix the distributed alphabet $\langle \mathbb{P}, dom : \Sigma \to (2^{\mathbb{P}} \setminus \emptyset)\rangle$, the leaf process $r$ and its parent $q$, and a Zielonka automaton with a correctness condition $\mathcal{A} = \langle \{S_p\}_{p\in\mathbb{P}}, s_{in}, \{\delta_a\}_{a\in\Sigma}, \{Corr_p\}_{p\in\mathbb{P}}\rangle$.

The first step in proving Theorem 7 is to simplify the problem. First, we can restrict to controllers of a special form called covering controllers. Next, we show that the component of $\mathcal{A}$ to be eliminated, that is $\mathcal{A}_r$, can be assumed to have a particular property ($r$-short). After these preparatory results we will be able to present the reduction of $\mathcal{A}$ to $\mathcal{A}^\triangledown$ (Section 3.3).

## 3.1 Covering controllers

The notion of a covering controller will simplify the presentation because it will allow us to focus on the runs of the controller instead of a product of the plant and the controller.

**Definition 8 (Covering controller)** Let $\mathcal{C}$ be a Zielonka automaton over the same alphabet as $\mathcal{A}$; let $C_p$ be the set of states of process $p$ in $\mathcal{C}$. Automaton

$\mathcal{C}$ is a *covering controller* for $\mathcal{A}$ if there is a function $\pi : \{C_p\}_{p\in\mathbb{P}} \to \{S_p\}_{p\in\mathbb{P}}$, mapping each $C_p$ to $S_p$ and satisfying two conditions: (i) if $c_{dom(b)} \xrightarrow{b} c'_{dom(b)}$ then $\pi(c_{dom(b)}) \xrightarrow{b} \pi(c'_{dom(b)})$; (ii) for every uncontrollable action $a$: if $a$ is enabled from $\pi(c_{dom(a)})$ then it is also enabled from $c_{dom(a)}$.

**Remark 9** Strictly speaking, a covering controller $\mathcal{C}$ may not be a controller since we do not require that every uncontrollable action is enabled in every state, but only those actions that are enabled in $\mathcal{A}$. From $\mathcal{C}$ one can get a controller $\hat{\mathcal{C}}$ by adding self-loops for all missing uncontrollable transitions.

Notice that thanks to the projection $\pi$, a covering controller can inherit the correctness condition of $\mathcal{A}$. Moreover, the sequences labeling the maximal runs of $\mathcal{C}$, $\mathcal{A}\times\mathcal{C}$ and $\mathcal{A}\times\hat{\mathcal{C}}$ are the same.

**Lemma 10** There is a correct controller for $\mathcal{A}$ if and only if there is a covering controller $\mathcal{C}$ for $\mathcal{A}$ such that all the maximal runs of $\mathcal{C}$ satisfy the inherited correctness condition.

**Proof**
If $\mathcal{C}$ is a covering controller for $\mathcal{A}$ such that all its maximal runs satisfy the inherited correctness condition then $\hat{\mathcal{C}}$ is a correct controller for $\mathcal{A}$. Conversely, if $\mathcal{C}$ is a correct controller for $\mathcal{A}$ then $\mathcal{A}\times\mathcal{C}$ is a covering controller where all maximal runs satisfy the inherited correctness condition. □

We will refer to a covering controller with the property that all its maximal runs satisfy the inherited correctness condition, as *correct covering controller*.

## 3.2 Short automata

In this section we justify our restriction to plants $\mathcal{A}$ where the $r$-component $\mathcal{A}_r$ is short (see Definition 11 below). Recall that we have assumed that all controllable actions are local and that we consider a tree architecture with a leaf process $r$ and its parent $q$.

**Definition 11 ($r$-short)** Automaton $\mathcal{A}$ is *$r$-short* if there is a bound on the number of actions that $r$ can perform without doing a communication with $q$.

**Theorem 12** *For every automaton $\mathcal{A}$, we can construct an $r$-short automaton $\mathcal{A}^{\mathbb{S}}$ such that there is a correct controller for $\mathcal{A}$ iff there is one for $\mathcal{A}^{\mathbb{S}}$.*

The rest of this subsection is devoted to the proof of the above theorem. Theorem 12 bears some resemblance with the fact that every parity game can be transformed into a finite game: when a loop is closed the winner is decided looking at the ranks on the loop. This construction would do if $r$ had no interaction with $q$. Possible interactions with $q$ make the construction more involved. Moreover, need to prove existence of some kind of memoryless strategies for distributed controllers.

Observe that we can make two simplifying assumptions. First, we assume that the correctness condition on $r$ is a parity condition. That is, it is given by a rank function $\Omega_r : S_r \to \mathbb{N}$ and the set of terminal states $T_r$. We can assume this since every regular language of infinite sequences can be recognized by a

deterministic parity automaton. The second simplification is to assume that the automaton $\mathcal{A}$ is *$r$-aware* with respect to the parity condition on $r$. This means that the state of $r$ determines the biggest rank that has been seen since the last communication of $r$ with $q$. It is easy to transform an automaton to an $r$-aware one.

Recall that if $\mathcal{C}$ is a covering controller for $\mathcal{A}$ (cf. Definition 8) then there is a function $\pi : \{C_p\}_{p\in\mathbb{P}} \to \{S_p\}_{p\in\mathbb{P}}$, mapping each $C_p$ to $S_p$ and respecting the transition relation: if $c_{dom(b)} \xrightarrow{b} c'_{dom(b)}$ then $\pi(c_{dom(b)}) \xrightarrow{b} \pi(c'_{dom(b)})$.

**Definition 13 ($r$-memoryless controller)** A covering controller $\mathcal{C}$ for $\mathcal{A}$ is *$r$-memoryless* when for every pair of states $c_r \neq c'_r$ of $\mathcal{C}_r$: if there is a path on local $r$-actions from $c_r$ to $c'_r$ then $\pi(c_r) \neq \pi(c'_r)$.

Intuitively, a controller can be seen as a strategy, and $r$-memoryless means that it does not allow the controlled automaton to go twice through the same $r$-state between two consecutive communication actions of $r$ and $q$.

**Lemma 14** Fix an $r$-aware automaton $\mathcal{A}$ with a parity correctness condition for process $r$. If there is a correct controller for $\mathcal{A}$ then there is also one that is covering and $r$-memoryless.

The proof of Lemma 14 uses the notion of signatures, that is classical in 2-player parity games, for defining a $r$-memoryless controller $\mathcal{C}^m$ from $\mathcal{C}$. The idea is to use representative states of $C_r$, defined in each strongly connected component according to a given signature and covering function $\pi$.

By Lemma 10 we can assume that we have a covering controller for $\mathcal{A}$. Let us fix an arbitrary linear order on the set $C_r$ of states of the automaton $\mathcal{C}_r$. Let $\mathcal{C}_r^{\downarrow loc}$ denote the graph obtained from $\mathcal{C}_r$ by taking $C_r$ as set of vertices and the transitions on local $r$-actions as edges. Since $\mathcal{C}$ is a covering controller, every sequence of actions in $\mathcal{C}$ can be performed in the controlled plant. Since $\mathcal{C}$ is correct for $\mathcal{A}$, every infinite sequence of local $r$-actions in the controlled plant satisfies the parity condition. We can lift this parity condition directly to $\mathcal{C}$ thanks to the fact that $\mathcal{C}$ is covering. We obtain that every infinite path in $\mathcal{C}_r^{\downarrow loc}$ satisfies the parity condition.

Before proceeding it will be convenient to recall some facts about parity games, in particular the notion of signature (or progress measure) [16]. We consider $\mathcal{C}_r^{\downarrow loc}$ as a parity game. Suppose that it uses priorities from $\{1, \ldots, d\}$. A signature is a $d$-tuple of natural numbers, that is, an element of $\mathbb{N}^d$. We will be interested in assignments of signatures to states of $\mathcal{C}_r^{\downarrow loc}$, that is in functions $sig : C_r \to \mathbb{N}^d$. Signatures are ordered lexicographically. We write $sig(c) \geq sig(c')$ if the signature assigned to $c$ is lexicographically bigger or equal to that of $c'$. For $i \in \{1, \ldots, d\}$ we write $sig(c) \geq_i sig(c')$ if the signature of $c$ truncated to the first $i$ positions is lexicographically bigger or equal to the signature of $c'$ truncated to the first $i$ positions. For a fixed assignment of signatures $sig$ and two states $c$, $c'$ of $\mathcal{C}_r$ we write $c \rhd_{sig} c'$ if

$$sig(c) \geq_{\Omega(c)} sig(c') \quad \text{and the inequality is strict if } \Omega(c) \text{ is odd.}$$

We say that an assignment of signatures $sig : C_r \to \mathbb{N}^d$ is *consistent* if for every edge $(c, c')$ of $\mathcal{C}_r^{\downarrow loc}$ we have $c \rhd_{sig} c'$. We now recall a fact that holds for every finite parity game, but we specialize them to $\mathcal{C}_r^{\downarrow loc}$.

**Fact.** Every path of $\mathcal{C}_r^{\downarrow loc}$ satisfies the parity condition iff there is a consistent assignment of signatures to states of $\mathcal{C}_r^{\downarrow loc}$.

After these preparations we can define for every state $c_r$ of $C_r$ its representative state in $C_r$, denoted $rep(c_r)$, as the unique state $c'_r$ satisfying the following conditions:

1. $\pi(c_r) = \pi(c'_r)$ and $c'_r$ is reachable from $c_r$;
2. for every $c''_r$ with $\pi(c_r) = \pi(c''_r)$: if $c''_r$ is reachable from $c'_r$ then it belongs to the same SCC as $c'_r$;
3. among all states satisfying points (1) and (2) consider those with the smallest signature; if there is more than one such state then pick the state that is the smallest in our fixed arbitrary ordering.

**Remark 15** For every $c'_r$ reachable in $\mathcal{C}_r^{\downarrow loc}$ from $rep(c_r)$: if $\pi(c'_r) = \pi(c_r)$ then $rep(c'_r) = rep(c_r)$. Indeed, by conditions (1) and (2) above $rep(c'_r)$ and $rep(c_r)$ must be in the same SCC. But then, the representative is uniquely determined by signature and ordering.

We define now $\mathcal{C}_r^m$ from $\mathcal{C}_r$ by redirecting every transition on a local $r$-action to representatives: if the transition goes to a state $c_r$ we make it go to $rep(c_r)$. Of course, $\mathcal{C}^m$ is still covering and the above remark implies that it is $r$-memoryless.

**Remark 16** If we have a transition $c_r \xrightarrow{b} c'_r$ in $\mathcal{C}_r^m$ then there is a sequence $z \in \Sigma_r^{loc}$ of local $r$-actions and some state $c''_r$ such that $c_r \xrightarrow{b} c''_r \xrightarrow{z} c'_r$ in $\mathcal{C}_r$.

Remark 16 allows to map paths in $\mathcal{C}_r^m$ into paths in $\mathcal{C}_r$. Consider a state $c_1$ of $\mathcal{C}_r^m$ and a finite sequence $x \in (\Sigma_r^{loc})^*$ such that $x = b_1 \cdots b_k$ labels some path from $c_1$ in $\mathcal{C}_r^m$, say $u = c_1 \xrightarrow{b_1} c_2 \xrightarrow{b_2} c_3 \cdots \xrightarrow{b_k} c_{k+1}$. Remark 16 gives us a sequence $rep^{-1}(c_1, x) = b_1 z_1 b_2 z_2 \dots b_k z_k$, and a corresponding path in $\mathcal{C}_r$: $c_1 \xrightarrow{b_1 z_1} c_2 \xrightarrow{b_2 z_2} c_3 \dots \xrightarrow{b_k z_k} c_{k+1}$, for some $z_i \in (\Sigma_r^{loc})^*$. In particular the two paths end in the same state. Of course $rep^{-1}(c_1, x)$ is defined similarly for infinite sequences $x$.

**Proof of Lemma 14.** We are ready to show that $\mathcal{C}^m$ obtained from $\mathcal{C}$ by replacing $\mathcal{C}_r$ with $\mathcal{C}_r^m$ satisfies the parity condition. For this take a maximal run and suppose towards a contradiction that it does not satisfy the parity condition.

If on this run there are infinitely many communications between $q$ and $r$ then there is an equivalent run whose labeling has the form:

$$u = y_0 x_0 a_1 y_1 x_1 a_2 \dots \tag{1}$$

where $a_i \in \Sigma_q \cap \Sigma_r$, $x_i \in (\Sigma_r^{loc})^*$, and $y_i \in (\Sigma \setminus \Sigma_r)^*$. Here two runs are equivalent means that the projections of the two runs on every process are identical. In particular, if two runs are equivalent and one of them satisfies the correctness condition then so does the other.

Let $c_r^i = state_r^{\mathcal{C}^m}(y_0 x_0 a_1 \cdots a_i)$ be the state of $\mathcal{C}_r^m$ reached on the prefix of $u$ up to $a_i$. Let $x'_i = rep^{-1}(c_r^i, x_i)$. We get that the sequence

$$u' = y_0 x'_0 a_1 y_1 x'_1 a_2 \dots \tag{2}$$

is a labeling of a maximal run in $\mathcal{C}$. The projections on processes other than $r$ are the same for $u$ and $u'$. It remains to see if the parity condition on $r$ is satisfied. We have $c_r^i \xrightarrow{x_i} c_r^{i+1}$ in $\mathcal{C}_r^m$ and $c_r^i \xrightarrow{x_i'} c_r^{i+1}$ in $\mathcal{C}_r$. Since we lifted priorities to $\mathcal{C}$ and $\mathcal{C}^m$ (being both covering), the $r$-awareness of $\mathcal{A}$ lifts to $\mathcal{C}$ and $\mathcal{C}^m$, so the same maximal rank is seen when reading $x_i$ and $x_i'$. This shows that the parity condition on $r$ is satisfied on the run of $\mathcal{C}^m$ on $u$, since it is satisfied by the run of $\mathcal{C}$ on $u'$.

Consider now a maximal run with finitely many communications between $q$ and $r$. There is an equivalent one labeled by a sequence of the form:

$$u = y_0 x_0 a_1 y_1 x_1 a_2 \cdots a_k y_k x_k \tag{3}$$

where $y_k$ and $x_k$ are potentially infinite. Since we have only modified the $r$-component of the controller, it must be $x_k$ that does not satisfy the parity condition on $r$.

Suppose first that $x_k = b_1 b_2 \cdots$ is infinite. Take the run $c_1 \xrightarrow{b_1} c_2 \xrightarrow{b_2} c_3 \ldots$ in $\mathcal{C}_r^m$, where $c_1 = state_r^{\mathcal{C}^m}(y_0 x_0 a_1 \cdots a_k)$. We have a run $c_1 \xrightarrow{b_1} c_2' \xrightarrow{x_2} c_2 \xrightarrow{b_2} c_3' \xrightarrow{x_3} c_3 \cdots$ in $\mathcal{C}_r$, where $rep(c_i') = c_i$ and $x_i \in (\Sigma_r^{loc})^*$ is the accessibility path, as given by Remark 16. We have $c_i \rhd_{sig} c_{i+1}'$ for all $i = 1, 2, \ldots$ because there is an edge from $c_i$ to $c_{i+1}'$ in $\mathcal{C}_r^{\downarrow loc}$. Recall that $c_i = rep(c_i')$. The definition of representatives implies that either $sig(c_i') \geq sig(c_i)$ or $c_i$ is in a strictly lower SCC than $c_i'$. Since lowering a component can happen only finitely many times we have $sig(c_i') \geq sig(c_i)$ for all $i$ bigger than some $n$. We get $c_i \rhd_{sig} c_{i+1}$ for $i > n$ which implies that $x_k$ satisfies the parity condition. A contradiction.

If $x_k$ is finite then we define the sequence $u' = y_0 x_0' a_1 y_1 x_1' a_2 \cdots a_k y_k x_k'$ in $L(\mathcal{C})$, as in the first case. Since $u$ was maximal in $\mathcal{C}^m$, we have that $u'$ is maximal in $\mathcal{C}$ (if $r$ can do an action in $u'$, the same can be done in $u$, since $u$ and $u'$ end in the same state). Thus the $r$-state reached in $\mathcal{A}$ by $u$ belongs to $T_r$, since this holds already for $u'$. We get again a contradiction.

We will use Lemma 14 to reduce the control problem to that for $r$-short automata.

Given $\mathcal{A}$ we define a $r$-short automaton $\mathcal{A}^{\circledS}$. All its components will be the same but for the component $r$. The states $S_r^{\circledS}$ of $r$ will be sequences $w \in S_r^+$ of states of $\mathcal{A}_r$ without repetitions, plus two new states $\top, \bot$. For a local transition $s_r' \xrightarrow{b} s_r''$ in $\mathcal{A}_r$ we have in $\mathcal{A}_r^{\circledS}$ transitions:

$$\begin{aligned}
&ws_r' \xrightarrow{b} ws_r's_r'' && \text{if } ws_r's_r'' \text{ a sequence without repetitions}\\
&ws_r' \xrightarrow{b} \top && \text{if } s_r'' \text{ appears in } w \text{ and the resulting loop is even}\\
&ws_r' \xrightarrow{b} \bot && \text{if } s_r'' \text{ appears in } w \text{ and the resulting loop is odd}
\end{aligned}$$

There are also communication transitions between $q$ and $r$:

$$(s_q, ws_r') \xrightarrow{b} (s_q', s_r'') \qquad \text{if } (s_q, s_r') \xrightarrow{b} (s_q', s_r'') \text{ in } \mathcal{A}$$

Notice that $w$ disappears in communication transitions. The parity condition for $\mathcal{A}^{\circledS}$ is also rather straightforward: it is the same for the components other than $r$, and for $\mathcal{A}_r$ it is

- $\Omega^{\circledS}(ws_r) = \Omega(s_r)$,

- $T_r^{\mathbb{S}} = \{\top\} \cup \{ws_r : s_r \in T_r\}$.

**Proof of Theorem 12:** Consider the implication from left to right. Let $\mathcal{C}$ be a correct covering controller for $\mathcal{A}$. By Lemma 14 we can assume that it is $r$-memoryless. We show that $\mathcal{C}$ is also a covering correct controller for $\mathcal{A}^{\mathbb{S}}$. We will concentrate on correctness, since the covering part follows by examination of the definitions.

Let us take some maximal run $run^{\mathbb{S}}(u)$ of $\mathcal{A}^{\mathbb{S}} \times \mathcal{C}$, and suppose by contradiction that it does not satisfy the parity condition of $\mathcal{A}^{\mathbb{S}}$. By definition $run(u)$ is a run of $\mathcal{A} \times \mathcal{C}$, but it may not be maximal. We have by construction of $\mathcal{A}^{\mathbb{S}}$ that $state_p^{\mathbb{S}}(u) = state_p(u)$ for $p \neq r$ and that $state_r(u)$ is the last element of $state_r^{\mathbb{S}}(u)$. (Recall that $state_p^{\mathbb{S}}(u), state_p(u)$ denote the state reached on $u$ by $\mathcal{A}^{\mathbb{S}} \times \mathcal{C}$ and $\mathcal{A} \times \mathcal{C}$, resp.)

Suppose that $run(u)$ is not a maximal run of $\mathcal{A} \times \mathcal{C}$. We will extend it to a maximal run $run(\overline{u})$. If $run(u)$ ended in $\bot$ in the $r$-component of $\mathcal{A}^{\mathbb{S}}$ then we could extend $run(u)$ to a run of $\mathcal{A} \times \mathcal{C}$ not satisfying the parity condition (here we use that $\mathcal{C}$ is memoryless, so the odd loop in $\mathcal{A}$ exists also into one in $\mathcal{A} \times \mathcal{C}$). So the only other possibility is that $run(u)$ ends in $\top$. In this case it is possible to extend $run(u)$ to a complete run of $\mathcal{A} \times \mathcal{C}$ by adding the even loop in the $r$-component. This makes $r$ satisfy the parity condition. Let $run(\overline{u})$ be the resulting run.

Now observe that if a parity condition for some process $p \neq r$ is violated on $u$ then on $\overline{u}$ the same condition is violated. If it is violated on $r$ then the only remaining possibility is that there are finitely many $r$-actions in $u$, and the state reached on $u$ is $ws_r$ with $s_r \notin T_r$. But then $u$ is a maximal run of $\mathcal{A} \times \mathcal{C}$ and is not well terminated on $r$ either, a contradiction.

For implication from right to left we take a covering controller $\mathcal{C}^{\mathbb{S}}$ for $\mathcal{A}^{\mathbb{S}}$ and construct a controller $\mathcal{C}$ for $\mathcal{A}$. The controller $\mathcal{C}$ will be obtained by modifying the $r$-component of $\mathcal{C}^{\mathbb{S}}$. The states of $\mathcal{C}_r$ will be sequences of states of $\mathcal{C}_r^{\mathbb{S}}$. They will be of bounded length. We will have that if $c_1^{\mathbb{S}} \cdots c_k^{\mathbb{S}}$ is a state of $\mathcal{C}_r$ then $\pi^{\mathbb{S}}(c_k^{\mathbb{S}})$ is the state $s_1 \cdots s_k$ of $\mathcal{A}_r^{\mathbb{S}}$, where $\pi^{\mathbb{S}}(c_j^{\mathbb{S}}) = s_1 \cdots s_j$, for $j = 1, \ldots, k$. Moreover, we define $\pi(c_1^{\mathbb{S}} \cdots c_k^{\mathbb{S}}) = s_k$. The transitions of $\mathcal{C}_r$ are

- $wc^{\mathbb{S}} \xrightarrow{b} wc^{\mathbb{S}}d^{\mathbb{S}}$ if $c^{\mathbb{S}} \xrightarrow{b} d^{\mathbb{S}}$ in $\mathcal{C}_r^{\mathbb{S}}$ and $\pi^{\mathbb{S}}(d^{\mathbb{S}}) \neq \top$.
- $c_1^{\mathbb{S}} \cdots c_k^{\mathbb{S}} \xrightarrow{b} c_1^{\mathbb{S}} \cdots c_j^{\mathbb{S}}$ if $c_k^{\mathbb{S}} \xrightarrow{b} c^{\mathbb{S}}$ in $\mathcal{C}_r^{\mathbb{S}}$, $\pi^{\mathbb{S}}(c^{\mathbb{S}}) = \top$ and $j$ is such that $\pi^{\mathbb{S}}(c_k^{\mathbb{S}})$ is $s_1 \cdots s_k$ with $s_k \xrightarrow{b} s_j$ in $\mathcal{A}$.

Notice that since $\mathcal{C}^{\mathbb{S}}$ satisfies the parity condition $\bot$ cannot be reached.

**Remark 17** The construction of $\mathcal{C}$ guarantees that every sequence of local $r$-actions $x$ of $\mathcal{C}$ has a corresponding (possibly shorter) sequence $x'$ of $\mathcal{C}^{\mathbb{S}}$. If the sequence in $\mathcal{C}^{\mathbb{S}}$ starts in $s_1$ and finishes in $s_2$ then the sequence in $\mathcal{C}$ starts also in $s_1$, but now considered as a sequence of length 1, and finishes in a sequence ending in $s_2$. Since $\mathcal{A}$ is $r$-aware and $\mathcal{C}$ and $\mathcal{C}^{\mathbb{S}}$ are both covering, this means that the maximal rank seen on both sequences is the same.

We need to show that all maximal runs of $\mathcal{A} \times \mathcal{C}$ satisfy the parity condition. For contradiction suppose that $run(u)$ does not.

If there are infinitely many communications between $q$ and $r$ on $u$ then we write it as

$$u = y_0 x_0 a_1 y_1 x_1 a_2 \dots \tag{4}$$

where $a_i \in \Sigma_q \cap \Sigma_r$, $x_i \in (\Sigma_r^{loc})^*$, and $y_i \in (\Sigma \setminus \Sigma_r)^*$. Now for every $x_i$, Observation 1 gives $x_i'$ so that the maximal ranks on $x_i$ and $x_i'$ are the same, so for

$$u' = y_0 x_0' a_1 y_1 x_1' a_2 \dots \tag{5}$$

$run(u')$ is a maximal run of $\mathcal{C}^{\circledS}$. This gives a run violating the parity condition of $\mathcal{C}^{\circledS}$.

If there are finitely many communications between $q$ and $r$ on $u$ then we write it as

$$u = y_0 x_0 a_1 y_1 x_1 a_2 \dots a_k y_k x_k \tag{6}$$

where $y_k$ and $x_k$ are potentially infinite. The only complicated case is when $x_k$ is infinite. We need to show that the run of $\mathcal{C}_r$ on $x_k$ satisfies the parity condition. Recall that the states of $\mathcal{C}_r$ are sequences of states of $\mathcal{C}_r^{\circledS}$. Moreover the length of this sequences is bounded. Take the shortest sequence appearing infinitely often in $x_k$. The biggest rank seen between consecutive appearances of this sequence is even, since the path can be decomposed into (several) even loops of $\mathcal{A}$.

**Corollary 18** In the $r$-short plant $\mathcal{A}^{\circledS}$ the $r$-controller may be chosen memoryless since there are no infinite local $r$-plays.

**Remark 19** We claim that the complexity of the reduction from $\mathcal{A}$ to $\mathcal{A}^{\triangledown}$ is polynomial in the size of $\mathcal{A}_q$ and simply exponential in the size of $\mathcal{A}_r$. The reason is as follows. States of $\mathcal{A}_r^{\circledS}$ are simple paths (i.e., without repetition of states) of $\mathcal{A}_r$. When going from $\mathcal{A}$ to $\mathcal{A}^{\triangledown}$, the states of $\mathcal{A}_q^{\triangledown}$ contain $r$-local strategies $f : (\Sigma_r^{loc})^* \to \Sigma_r^{sys}$. Putting things together, in $f$ we deal with paths of $\mathcal{A}_r^{\circledS}$ (mapping them to $\Sigma_r^{sys}$). But the latter can be written more succinctly as paths of $\mathcal{A}_r$ without repetitions.

## 3.3 The reduced automaton $\mathcal{A}^{\triangledown}$

Equipped with the notions of covering controller and $r$-short strategy we can now present the construction of the reduced automaton $\mathcal{A}^{\triangledown}$. We suppose that $\mathcal{A} = \langle \{S_p\}_{p \in \mathbb{P}^{\triangledown}}, s_{in}, \{\delta_a\}_{a \in \Sigma} \rangle$ is $r$-short and we define now the reduced automaton $\mathcal{A}^{\triangledown}$ that results by eliminating process $r$ (cf. Figure 1). Let $\mathbb{P}^{\triangledown} = \mathbb{P} \setminus \{r\}$. We construct $\mathcal{A}^{\triangledown} = \langle \{S_p^{\triangledown}\}_{p \in \mathbb{P}^{\triangledown}}, s_{in}^{\triangledown}, \{\delta_a^{\triangledown}\}_{a \in \Sigma^{\triangledown}} \rangle$ where the components are defined below.

All the processes $p \neq q$ of $\mathcal{A}^{\triangledown}$ will be the same as in $\mathcal{A}$. This means: $S_p^{\triangledown} = S_p$, and $\Sigma_p^{\triangledown} = \Sigma_p$. Moreover, all transitions $\delta_a$ with $dom(a) \cap \{q, r\} = \emptyset$ are as in $\mathcal{A}$. Finally, in $\mathcal{A}^{\triangledown}$ the correctness condition of $p \neq q$ is the same as in $\mathcal{A}$.

Before defining process $q$ in $\mathcal{A}^{\triangledown}$ let us introduce the notion of $r$-local strategy. An *$r$-local strategy from a state $s_r \in S_r$* is a partial function $f : (\Sigma_r^{loc})^* \to \Sigma_r^{sys}$ mapping sequences from $\Sigma_r^{loc}$ to actions from $\Sigma_r^{sys}$, such that if $f(v) = a$ then $s_r \xrightarrow{va}$ in $\mathcal{A}_r$. Observe that since the automaton $\mathcal{A}$ is $r$-short, the domain of $f$ is finite.

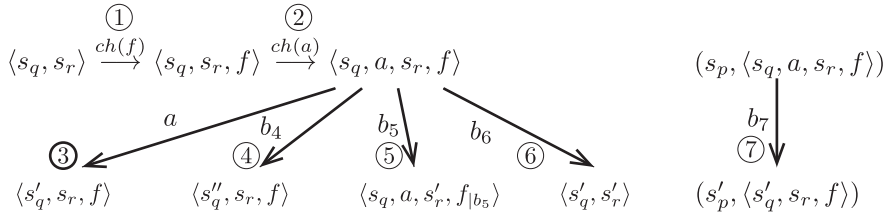


Figure 2: Transitions of $\mathcal{A}^\nabla$

Given an $r$-local strategy $f$ from $s_r$, a local action $a \in \Sigma_r^{loc}$ is *allowed by* $f$ if $f(\epsilon) = a$, or $a$ is uncontrollable. For $a$ allowed by $f$ we denote by $f_{|a}$ the $r$-local strategy defined by $f_{|a}(v) = f(av)$; this is a strategy from $s'_r$, where $s_r \xrightarrow{a} s'_r$.

The states of process $q$ in $\mathcal{A}^\nabla$ are of one of the following types:

$$\langle s_q, s_r\rangle\,, \quad \langle s_q, s_r, f\rangle\,, \quad \langle s_q, a, s_r, f\rangle\,,$$

where $s_q \in S_q, s_r \in S_r$, $f$ is a $r$-local strategy from $s_r$, and $a \in \Sigma_q^{loc}$. The new initial state for $q$ is $\langle (s_{in})_q, (s_{in})_r\rangle$. Recall that that since $\mathcal{A}$ is $r$-short, any $r$-local strategy in $\mathcal{A}_r$ is necessarily finite, so $S_q^\nabla$ is a finite set. Recall also that controllable actions are local.

The transitions of $\mathcal{A}_q^\nabla$ are presented in Figure 2. Transition ① chooses an $r$-local strategy $f$. It is followed by transition ② that declares a controllable action $a \in \Sigma_q^{sys}$ that is enabled from $s_q$. Transition ③ executes the chosen action $a$; we require $s_q \xrightarrow{a} s'_q$ in $\mathcal{A}_q$. Transition ④ executes an uncontrollable local action $b_4 \in \Sigma_q^{env}$; provided $s_q \xrightarrow{b_4} s''_q$ in $\mathcal{A}_q$. Transition ⑤ executes a local action $b_5 \in \Sigma_r^{loc}$, provided that $b_5$ is allowed by $f$ and $s_r \xrightarrow{b_5} s'_r$. Transition ⑥ simulates a synchronization $b_6$ between $q$ and $r$; provided $(s_q, s_r) \xrightarrow{b_6} (s'_q, s'_r)$ in $\mathcal{A}$. Finally, transition ⑦ simulates a synchronization between $q$ and $p \neq r$. An example of a simulation of $\mathcal{A}_q$ and $\mathcal{A}_r$ by $\mathcal{A}_q^\nabla$ is presented in Figure 3. The numbers below transitions refer to the corresponding cases from the definition.

To summarize, in $\Sigma_q^\nabla$ we have all actions of $\Sigma_r$ and $\Sigma_q$, but they become uncontrollable. All the new actions of process $q$ in plant $\mathcal{A}^\nabla$ are *controllable*:

- action $ch(f) \in \Sigma^{sys}$, for every local $r$-strategy $f$,
- action $ch(a)$, for every $a \in \Sigma_q^{sys}$.

The correctness condition for process $q$ in $\mathcal{A}^\nabla$ is:

1. The correct infinite runs of $q$ in $\mathcal{A}^\nabla$ are those that have the projection on transitions of $\mathcal{A}_q$ correct with respect to $Corr_q$, and either: *(i)* the projection on transitions of $\mathcal{A}_r$ is infinite and correct with respect to $Corr_r$; or *(ii)* the projection on transitions of $\mathcal{A}_r$ is finite and for $f, s_r$ appearing in almost all states of $q$ of the run we have that from $s_r$ all sequences respecting strategy $f$ end in a state from $T_r$.
2. $T_q^\nabla$ contains states $\langle s_q, s_r, f\rangle$ such that $s_q \in T_q$, and $s_r \in T_r$.

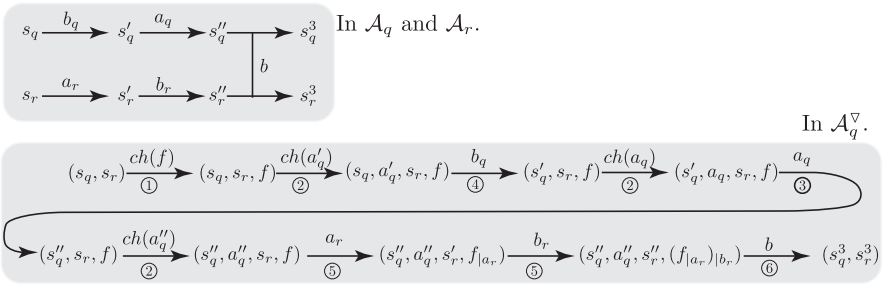


Figure 3: Simulation of $\mathcal{A}_q$ and $\mathcal{A}_r$ by $\mathcal{A}_q^\nabla$.

Item 1($ii$) in the definition above captures the case where $q$ progresses alone till infinity and blocks $r$, even though $r$ could reach a terminal state in a couple of moves. Clearly, item 1 can be expressed as an $\omega$-regular condition. The definition of correctness condition is one of the principal places where the $r$-short assumption is used. Without this assumption we would need to cope with the situation where we have an infinite execution of $\mathcal{A}_q$, and at the same time an infinite execution of $\mathcal{A}_r$ that do not communicate with each other. In this case $\mathcal{A}_q^\nabla$ would need to fairly simulate both executions in some way.

The reduction is rather delicate since in concurrent systems there are many different interactions that can happen. For example, we need to schedule actions of process $q$, using $ch(a)$ actions, before the actions of process $r$. The reason is the following. First, we need to make all $r$-actions uncontrollable, so that the environment could choose any play respecting the chosen $r$-local strategy. Now, if we allowed controllable $q$-actions to be eligible at the same time as $r$-actions, then the control strategy for automaton $\mathcal{A}^\nabla$ would be to propose nothing and force the environment to play the $r$-actions. This would allow the controller of $\mathcal{A}^\nabla$ to force the advancement of the simulation of $r$ and get information that is impossible to obtain by the controller of $\mathcal{A}$.

Together with Theorem 12, the theorem below implies our main Theorem 7.

**Theorem 20** *For every $r$-short Zielonka automaton $\mathcal{A}$ and every local, $\omega$-regular correctness conditions: there is a correct covering controller for $\mathcal{A}$ iff there is a correct covering controller for $\mathcal{A}^\nabla$. The size of $\mathcal{A}_q^\nabla$ is polynomial in the size of $\mathcal{A}_q$ and exponential in the size of $\mathcal{A}_r$.*

We end with some notations used in the following sections. For a Zielonka automaton $\mathcal{C}^\nabla$ over $(\Sigma^\nabla, loc)$ and $w \in (\Sigma^\nabla)^\infty$ we will write $run^\nabla(w)$ for the sequence of transitions of $\mathcal{C}^\nabla$ when reading $w$. For finite $w$ we will write $state(w)$ for the last state in $run^\nabla(w)$.

## 3.4 Proof of Theorem 20: from $\mathcal{C}$ to $\mathcal{C}^\nabla$

By Lemma 10 we can assume that we have a *correct covering controller* $\mathcal{C}$ for $\mathcal{A}$. We show how to construct a correct controller $\mathcal{C}^\nabla$ for $\mathcal{A}^\nabla$. This will give the left to right implication of Theorem 20.

**Remark 21** Some simple observations about $\mathcal{C}$.

1. We may assume that from every state of $C$ there is at most one transition on a local controllable action. If there were more than one, we could arbitrary remove one of them. This will reduce the number of maximal runs so the resulting controller with stay correct.

2. $\mathcal{C}$ determines for every state $c$ of $\mathcal{C}_r$ a local $r$-strategy $f$ from $\pi(c)$: if $c = c_0 \xrightarrow{a_1} c_1 \xrightarrow{a_2} \cdots \xrightarrow{a_k} c_k$, $\pi(c_i) = s_i$ and $a_i \in \Sigma_r^{loc}$ for all $i$, then $f(c_0 \cdots c_k) = a$ where $a \in \Sigma_r^{sys}$ is a (unique) controllable action possible from $c_k$. This strategy may have memory, but all the (local) plays respecting $f$ are of bounded length, assuming that $\mathcal{A}$ is $r$-short.

The components $\mathcal{C}_p^\nabla$ for $p \neq q$ are just $\mathcal{C}_p$, and the initial state is the same. The component $\mathcal{C}_q^\nabla$ is described below. Its states are of the form $(c_q, c_r)$, $(c_q, c_r, f)$ and $(c_q, a, c_r, f)$ with $c_q \in C_q$, $c_r \in C_r$, $a \in \Sigma_q^{sys}$, and local $r$-strategy $f$. Its initial state is $(c_q^0, c_r^0)$, with $c_q^0, c_r^0$ initial states of $\mathcal{C}_q, \mathcal{C}_r$.

The transitions of $\mathcal{C}_q^\nabla$ ensure the right choice of a local strategy and of a local action:

- Choice of $r$-strategy:
$$(c_q, c_r) \xrightarrow{ch(f)} (c_q, c_r, f)$$
where $f$ is the local $r$-strategy from $\pi(c_r)$ determined by $\mathcal{C}$ in state $c_r$.

- Choice of a (local) controllable $q$-action:
$$(c_q, c_r, f) \xrightarrow{ch(a)} (c_q, a, c_r, f)$$
For $a \in \Sigma_q^{sys}$ unique such that $c_q \xrightarrow{a} c_q'$, for some $c_q'$. If there is no such transition then we put some arbitrary fixed action $a_0 \in \Sigma_q^{sys}$.

The other transitions of $\mathcal{C}_q^\nabla$ are on uncontrollable actions, they just reflect the structure of $\mathcal{A}^\nabla$:

- Execution of the chosen controllable $q$-action:
$$(c_q, a, c_r, f) \xrightarrow{a} (c_q', c_r, f) \qquad \text{if } c_q \xrightarrow{a} c_q' \text{ in } \mathcal{C}_q.$$

- Execution of an uncontrollable local $q$-action:
$$(c_q, a, c_r, f) \xrightarrow{b} (c_q', c_r, f) \qquad \text{if } c_q \xrightarrow{b} c_q' \text{ in } \mathcal{C}_q, \text{ where } b \in \Sigma_q^{env} \cap \Sigma_q^{loc}.$$

- Communication between $q$ and $p \neq r$:
$$(c_p, (c_q, a, c_r, f)) \xrightarrow{b} (c_p', (c_q', c_r, f)) \quad \text{if } (c_p, c_q) \xrightarrow{b} (c_p', c_q') \text{ in } \mathcal{C}.$$

- Local move of $r$:
$$(c_q, a, c_r, f) \xrightarrow{b} (c_q, a, c_r', f) \quad \text{if } c_r \xrightarrow{b} c_r' \text{ in } \mathcal{C}_r, \text{ where } b \in \Sigma_r^{loc}.$$

- Communication between $q$ and $r$
$$(c_q, a, c_r, f) \xrightarrow{b} (c_q', c_r') \quad \text{if } (c_q, c_r) \xrightarrow{b} (c_q', c_r') \text{ in } \mathcal{C}.$$

**Lemma 22** If $C$ is a covering controller for $\mathcal{A}$ then $\mathcal{C}^\nabla$ is a covering controller for $\mathcal{A}^\nabla$. The covering function is

$$\begin{aligned}\pi^\nabla(c_q, c_r) =&(\pi(c_q), \pi(c_r))\\ \pi^\nabla(c_q, c_r, f) =&(\pi(c_q), \pi(c_r), f)\\ \pi^\nabla(c_q, a, c_r, f) =&(\pi(c_q), a, \pi(c_r), f)\ .\end{aligned}$$

For the correctness proof we will need one more definition:

**Definition 23 (*hide*)** For $w \in (\Sigma^\nabla)^\infty$ we let $hide(w) \in \Sigma^\infty$ be the sequence obtained by removing actions from $\Sigma^\nabla \setminus \Sigma$.

Observe that by construction of $\mathcal{C}^\nabla$ if $run^\nabla(w)$ is defined then in $w$ there can be at most two consecutive $q$-actions from $\Sigma^\nabla \setminus \Sigma$.

**Lemma 24** Let $w \in (\Sigma^\nabla)^*$. If $run^\nabla(w)$ is defined then so is $run(hide(w))$. Moreover, letting $c^\nabla = state^\nabla(w)$ and $c = state(hide(w))$, we have that (i) $c^\nabla_p = c_p$ for all $p \neq q, r$, and (ii) $c^\nabla_q$ is either $(c_q, c_r)$, or $(c_q, c_r, f)$, or $(c_q, a, c_r, f)$; where $a$ and $f$ are determined by $c_q$ and $c_r$ as follows:

- $a$ is the unique controllable $q$-action from $c_q$ in $\mathcal{C}$ (or $a_0$ if there is none).
- $f$ is the local $r$-strategy determined by $\mathcal{C}$ in $c_r$.

**Proof**
The proof is by induction on the length of $w$. It follows by direct examination of the rules. □

**Lemma 25** Assume that $w \in (\Sigma^\nabla)^\infty$. For every process $p \neq q$ we have $run^\nabla_p(w) = run_p(hide(w))$. Concerning $run^\nabla_q(w)$: if we project it on transitions of $\mathcal{C}_q$ we obtain $run_q(hide(w))$; if we project it on transitions of $\mathcal{C}_r$ we obtain $run_r(hide(w))$.

**Proof**
Directly from the previous lemma. □

**Lemma 26** If $\mathcal{C}$ is a correct covering controller for $\mathcal{A}$ then $\mathcal{C}^\nabla$ is a correct covering controller for $\mathcal{A}^\nabla$.

**Proof**
Since $\mathcal{C}$ is a correct covering controller we have that all maximal runs of $\mathcal{C}$ are correct w.r.t $\mathcal{A}$. By Lemma 22 we know that $\mathcal{C}^\nabla$ is a covering controller, so it is enough to show that all maximal runs of $\mathcal{C}^\nabla$ are correct w.r.t. $\mathcal{A}^\nabla$.

Take a maximal run in $\mathcal{C}^\nabla$, say on $w \in (\Sigma^\nabla)^\infty$. The first obstacle is that $run(hide(w))$ may be not maximal in $\mathcal{C}$. This can only happen when there are infinitely many $q$-actions in $w$, but only finitely many $r$-actions. Then we have $w = v_1 v_2$ and there are no $r$-actions in $v_2$. Let $state^\nabla_q(v_1) = (c_q, a, c_r, f)$. We have that $c_r$ and $f$ appear in all $state^\nabla_q(v_1 v')$, for every prefix $v'$ of $v_2$. The run $run(hide(w))$ is not maximal when there is at least some local action of $\mathcal{C}_r$ enabled in $c_r$. Let $x$ be a maximal sequence of local $r$-actions that is possible in $\mathcal{C}_r$ from state $c_r$. Since $\mathcal{A}$ is $r$-short, every such sequence is finite. Moreover we

choose $x$ in such a way that it brings $\mathcal{C}_r$ into a state not in $T_r$ (if it is possible). We get that $u = v_1 x v_2$ also defines a maximal run of $\mathcal{C}^\triangledown$, but now the run on $hide(u)$ is maximal in $\mathcal{C}$. Notice that $run^\triangledown(u)$ satisfies $Corr^\triangledown$ iff $run^\triangledown(w)$ does: the difference is the sequence $x$, and we have chosen, if possible, a losing sequence.

We need to show that the run of $\mathcal{A}^\triangledown$ on $u$ satisfies $Corr^\triangledown$ using the fact that the run on $hide(u)$ satisfies $Corr$. For $p \neq q$, Lemma 25 tells us that $run_p^\triangledown(u)$ is the same as $run_p(hide(u))$. Since $Corr_p^\triangledown$ and $Corr_p$ are the same, we are done.

It remains to consider $run_q^\triangledown(u)$. If there are finitely many $q$-actions in $u \in (\Sigma^\triangledown)^\infty$ then $u = u_1 u_2$ with no $q$-action in $u_2$. Consider $state_q^\triangledown(u_1) = (c_q, a, c_r, f)$. We have that $state_q(hide(u_1)) = c_q$ and $state_r(hide(u_1)) = c_r$. As there are no $q$-actions in $u_2$, and $run_q(u)$ satisfies $Corr_q$, we must have $\pi(c_q) \in T_q$ and $\pi(c_r) \in T_r$. This shows that $run_q^\triangledown(u)$ satisfies $Corr_q^\triangledown$.

If there are infinitely many $q$-actions in $u \in (\Sigma^\triangledown)^\infty$, we still have two cases. The first is when there are infinitely many actions from $\Sigma_r$ as well. Then $run_q^\triangledown(u)$ satisfies $Corr_q^\triangledown$ if the corresponding runs $run_q(hide(u))$ and $run_r(hide(u))$ satisfy $Corr_q$ and $Corr_r$, respectively. This is guaranteed by our assumption that $run(hide(u))$ satisfies $Corr$.

The last case is when in $u \in (\Sigma^\triangledown)^\infty$ we have infinitely many $q$-actions and only finitely many actions from $\Sigma_r$. Then $u = u_1 u_2$ with no actions from $\Sigma_r$ in $u_2$. We get $state_q^\triangledown(u_1) = (c_q, a, c_r, f)$ with both $c_r$, $f$ appearing in all the further states of the run. Since $run_r(hide(u))$ satisfies $Corr_r$, we have that $\pi(c_r) \in T_r$. But then, by the construction of $u$, there is no $\Sigma_r$-transition possible from $c_r$ (and neither from $\pi(c_r)$ in $\mathcal{A}_r^\triangledown$, since $\mathcal{C}^\triangledown$ is covering). This means that $run_q^\triangledown(u)$ satisfies $Corr_q^\triangledown$.

□

## 3.5 Proof of Theorem 20: from $\mathcal{D}^\triangledown$ to $\mathcal{D}$

This subsection gives the right-to-left direction of the proof. Given a correct controller $\mathcal{D}^\triangledown$ for $\mathcal{A}^\triangledown$, we show how to construct a correct controller $\mathcal{D}$ for $\mathcal{A}$. By Lemma 10 we can assume that $\mathcal{D}^\triangledown$ is covering.

The components $\mathcal{D}_p$ for $p \neq q, r$ are the same as in $\mathcal{D}^\triangledown$. So it remains to define $\mathcal{D}_q$ and $\mathcal{D}_r$. The states of $\mathcal{D}_q$ and $\mathcal{D}_r$ are obtained from states of $\mathcal{D}_q^\triangledown$. We need only certain states of $\mathcal{D}_q^\triangledown$, namely those $d_q$ whose projection $\pi^\triangledown(d_q)$ in $\mathcal{A}_q^\triangledown$ has four components, we call them *true states* of $\mathcal{D}_q^\triangledown$:

$$ts(\mathcal{D}_q^\triangledown) = \{d_q \in \mathcal{D}_q^\triangledown \mid \pi^\triangledown(d_q) \text{ is of the form } (s_q, a, s_r, f)\}.$$

Figure 4 presents an execution of $\mathcal{A}^\triangledown$ controlled by $\mathcal{D}^\triangledown$. We can see that $d_2$ is a true state, and $d_3$ is not.

The set of states of $\mathcal{D}_q$ is just $ts(\mathcal{D}_q^\triangledown)$, while the states of $\mathcal{D}_r$ are pairs $(d_q, x)$ where $d_q$ is a state from $ts(\mathcal{D}_q^\triangledown)$ and $x \in (\Sigma_r^{loc})^*$ is a sequence of local $r$-actions that is possible from $d_q$ in $\mathcal{D}^\triangledown$, in symbols $d_q \xrightarrow{x}$. We will argue later that such sequences are uniformly bounded. The initial state of $\mathcal{D}_q$ is the state $d_q^1$ reached from the initial state of $\mathcal{D}_q^\triangledown$ by the (unique) transitions of the form $ch(f_0), ch(a_0)$. The initial state of $\mathcal{D}_r$ is $(d_q^1, \varepsilon)$. The local transitions for $\mathcal{D}_r$ are $(d_q, x) \xrightarrow{b} (d_q, xb)$, for every $b \in \Sigma_r^{loc}$ and $d_q \xrightarrow{xb}$.

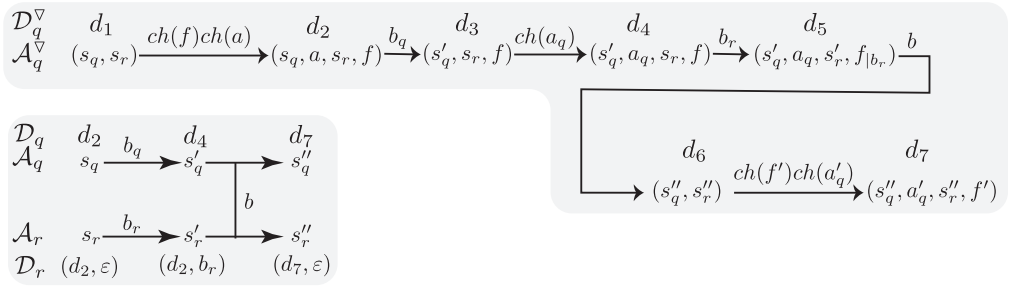
$\mathcal{D}_q^\nabla$ $d_1$ $ch(f)ch(a)$ $d_2$ $b_q$ $d_3$ $ch(a_q)$ $d_4$ $b_r$ $d_5$ $b$
$\mathcal{A}_q^\nabla$ $(s_q, s_r)$ $(s_q, a, s_r, f)$ $(s'_q, s_r, f)$ $(s'_q, a_q, s_r, f)$ $(s'_q, a_q, s'_r, f_{|b_r})$
$d_6$ $d_7$
$(s''_q, s''_r)$ $ch(f')ch(a'_q)$ $(s''_q, a'_q, s''_r, f')$
$\mathcal{D}_q$ $d_2$ $b_q$ $d_4$ $d_7$
$\mathcal{A}_q$ $s_q$ $s'_q$ $s''_q$
$b$
$\mathcal{A}_r$ $s_r$ $b_r$ $s'_r$ $s''_r$
$\mathcal{D}_r$ $(d_2, \varepsilon)$ $(d_2, b_r)$ $(d_7, \varepsilon)$

Figure 4: Decomposing controller $\mathcal{D}_q^\nabla$ into $\mathcal{D}_q$ and $\mathcal{D}_r$.

Before defining the transitions of $\mathcal{D}_q$ let us observe that if $d_q \in \mathcal{D}_q^\nabla$ is not in $ts(\mathcal{D}_q^\nabla)$ then only one controllable transition is possible from it. Indeed, as $\mathcal{D}^\nabla$ is a covering controller, if $\pi^\nabla(d_q)$ is of the form $(s_q, s_r)$ then there can be only an outgoing transition on a letter of the form $ch(f)$. Similarly, if $\pi^\nabla(d_q)$ is of the form $(s_q, s_r, f)$ then only a $ch(a)$ transition is possible. Since both $ch(f)$ and $ch(a)$ are controllable, we can assume that in $\mathcal{D}_q^\nabla$ there is no state with two outgoing transitions on a letter of this form. For a state $d_q \in \mathcal{D}_q^\nabla$ not in $ts(\mathcal{D}_q^\nabla)$ we will denote by $ts(d_q)$ the unique state of $ts(\mathcal{D}_q^\nabla)$ reachable from $d_q$ by one or two transitions of the kind $\xrightarrow{ch(f)}$ or $\xrightarrow{ch(a)}$, depending on the cases discussed above. Going back to Figure 4, we have $ts(d_3) = d_4$.

We now describe the $q$-actions possible in $\mathcal{D}$.

- Local $q$-action $b \in \Sigma_q^{loc}$: $d_q \xrightarrow{b} ts(d'_q)$ if $d_q \xrightarrow{b} d'_q$ in $\mathcal{D}_q^\nabla$. For example, this gives a transition $d_1 \xrightarrow{b_q} d_4$ in Figure 4.
- Communication $b \in \Sigma_q \cap \Sigma_p$ between $q$ and $p \neq r$: $(d_p, d_q) \xrightarrow{b} (d'_p, ts(d'_q))$ if $(d_p, d_q) \xrightarrow{b} (d'_p, d'_q)$ in $\mathcal{D}^\nabla$.
- Communication $b \in \Sigma_q \cap \Sigma_r$ of $q$ and $r$: $(d_q^1, (d_q^2, x)) \xrightarrow{b} (ts(d''_q), (ts(d''_q), \varepsilon))$ if $d_q^1 \xrightarrow{x} d'_q \xrightarrow{b} d''_q$ in $\mathcal{D}_q^\nabla$; observe that $\xrightarrow{x}$ is a sequence of transitions. For example, this gives a transition on $b$ in Figure 4.

In the last item the transition does not depend on $d_q^2$ since, informally, $d_q^1$ has been reached from $d_q^2$ by a sequence of actions independent of $r$. The condition $d_q^1 \xrightarrow{x} d'_q \xrightarrow{b} d''_q$ simulates the order of actions where all local $r$-actions come after the other actions of $q$, then we add a communication between $q$ and $r$.

The next lemma says that $\mathcal{D}$ is a covering controller for $\mathcal{A}$. Since $\mathcal{A}$ is assumed to be $r$-short, the lemma also gives a bound on the length of sequences in the states of $\mathcal{D}_r$.

**Lemma 27** If $\mathcal{D}^\nabla$ is a covering controller for $\mathcal{A}^\nabla$ then $\mathcal{D}$ is a covering controller for $\mathcal{A}$.

**Proof**
We need to define the projection function $\pi$ using the projection function $\pi^\nabla$. For $p \neq q, r$ set $\pi = \pi^\nabla$. For $\mathcal{D}_q$ we define $\pi(d_q) = s_q$ where $s_q$ is the state of $\mathcal{A}_q$

in $\pi^\triangledown(d_q)$. For $\mathcal{D}_r$ and its state $(d_q, x)$ we define $\pi(d_q, x) = s'_r$ where $d_q \xrightarrow{x} d'_q$ and $s'_r$ is the state of $\mathcal{A}_r$ in $\pi^\triangledown(d'_q)$.

We need to check that the transitions defined above preserve this projection function; namely for every process $p$: if $d_p \xrightarrow{b} d'_p$ in $\mathcal{D}_p$ then $\pi(d_p) \xrightarrow{b} \pi(d'_p)$ in $\mathcal{A}_p$; and similarly for communication actions. The statement is obvious if the move is in components other than $q$ or $r$. We are left with four cases:

- Local move of $q$, namely $d_q \xrightarrow{b} d'_q$. We have $d_q \xrightarrow{b} d''_1 \xrightarrow{ch(a')} d'_q$ in $\mathcal{D}^\triangledown$ for some $a'$, since $d'_q = ts(d''_q)$. By the fact that $\mathcal{D}^\triangledown$ covers $\mathcal{A}^\triangledown$ and the definition of moves of the latter automaton we have in $\mathcal{A}^\triangledown$:

$$\pi^\triangledown(d_q) = \langle s_q, a, s_r, f\rangle \xrightarrow{b} \langle s'_q, s_r, f\rangle \xrightarrow{ch(a')} \langle s'_q, a', s_r, f\rangle = \pi^\triangledown(d'_q)\,,$$

  and by definition of $\mathcal{A}^\triangledown$ we know that $s_q \xrightarrow{b} s'_q$ is in $\mathcal{A}$.

- Communication between $q$ and $p \neq r$ is similar.

- Local move of $r$: $(d_q, x) \xrightarrow{b} (d_q, xb)$. By definition we know that from $d_q$ it is possible to do in $\mathcal{D}^\triangledown$ the sequence of actions $xb$, that is $d_q \xrightarrow{x} d^1_q \xrightarrow{b} d^2_q$. We have $\pi^\triangledown(d_q) = (s_q, a, s_r, f)$, $\pi^\triangledown(d^1_q) = (s_q, a, s^1_r, f|_x)$ and $\pi^\triangledown(d^2_q) = (s_q, a, s^2_r, f|_{xb})$; since $xb$ is a sequence of local $r$-actions the other components do not change. We have $s^1_r \xrightarrow{b} s^2_r$ by definition of $\mathcal{A}^\triangledown_q$, and $\pi^\triangledown(d_q, x) = s^1_r$, $\pi^\triangledown(d_q, xb) = s^2_r$, as required.

- Communication between $q$ and $r$: $(d^1_q, (d^2_q, x)) \xrightarrow{b} (ts(d''_q), (ts(d''_q), \varepsilon))$. By definition this is possible only when $d^1_q \xrightarrow{x} d'_q \xrightarrow{b} d''_q$ in $\mathcal{D}^\triangledown_q$. Since $\mathcal{D}^\triangledown$ is covering we get the following sequence of transitions in $\mathcal{A}^\triangledown$:

$$\pi^\triangledown(d^1_q) = \langle s_q, a, s_r, f\rangle \xrightarrow{x} \langle s_q, a, s^1_r, f|_x\rangle \xrightarrow{b} \langle s'_q, s'_r\rangle \xrightarrow{ch(g)}$$
$$\langle s'_q, s'_r, g\rangle \xrightarrow{ch(a')} \langle s'_q, a', s'_r, g\rangle = \pi^\triangledown(ts(d''_q))$$

So we have $(s_q, s^1_r) \xrightarrow{b} (s'_q, s'_r)$ in $\mathcal{A}$ and $\pi(d^1_q) = s_q$, $\pi(ts(d''_q)) = s'_q$, $\pi(ts(d''_q), \varepsilon) = s'_r$. We claim that $\pi(d^2_q, x) = s^1_r$, and for this we need to observe a property of the runs of $\mathcal{D}$ (proved by induction on the length of the run). The intuition for the property below is that $d_q$ was reached from $d'_q$ by actions that do not involve $r$.

> **Property (*)** If from the initial state $\mathcal{D}$ can reach a global state with $d_q$ and $(d'_q, x)$ at the coordinates corresponding to $q$ and $r$, respectively, then the $s_r$- and $f$-components of the $\pi^\triangledown$ projections of $d_q$ and $d'_q$ are the same: $\pi^\triangledown(d_q) = (s_q, a, s_r, f)$ and $\pi^\triangledown(d'_q) = (s'_q, a', s_r, f)$, for some $s_q, s'_q, a, a', s_r, f$.

From Property (*) it follows that $\pi(d^2_q, \varepsilon) = s_r$, hence $\pi(d^2_q, x) = s^1_r$ since $s_r \xrightarrow{x} s^1_r$.

It remains to check the controllability condition for $\mathcal{D}$. For components other than $q$ and $r$ this is obvious. We have four cases to examine.

First, let us take a state $(d_q, x)$ of $\mathcal{D}_r$. Suppose that $\pi(d_q, x) \xrightarrow{b} s'_r$ is a local, uncontrollable transition in $\mathcal{A}_r$. We need to show that $(d_q, x) \xrightarrow{b} (d_q, xb)$ is possible in $\mathcal{D}_r$. Since $(d_q, x)$ is a state of $\mathcal{D}_r$ we have $d_q \xrightarrow{x} d'_q$ in $\mathcal{D}_q^\nabla$. Moreover, $\pi^\nabla(d'_q)$ is of the form $(s_q, a, s_r, f)$ and $\pi(d_q, x) = s_r$. We get that $(s_q, a, s_r, f) \xrightarrow{b} (s_q, s'_r, f|_b)$ exists in $\mathcal{A}_q^\nabla$. Since $\mathcal{D}^\nabla$ satisfies the controllability condition, in $\mathcal{D}_q^\nabla$ there must be a transition $d'_q \xrightarrow{b} d''_q$ for some $d''_q$. Hence, by definition, $(d_q, x) \xrightarrow{b} (d_q, xb)$ exists in $\mathcal{D}_r$.

For the next case we take a state $d_q$ of $\mathcal{D}_q$ and suppose that $\pi(d_q) \xrightarrow{b} s'_q$ is a local, uncontrollable transition in $\mathcal{A}_q$. We need to show that a $b$-transition is possible from $d_q$ in $\mathcal{D}_q$. We get $\pi^\nabla(d_q)$ is of the form $(s_q, a, s_r, f)$, and $\pi(d_q) = s_q$. This means that the transition $(s_q, a, s_r, f) \xrightarrow{b} (s'_q, s_r, f)$ is in $\mathcal{A}_q^\nabla$. Since $\mathcal{D}^\nabla$ is covering, we get $d_q \xrightarrow{b} d'_q$ for some $d'_q$ in $\mathcal{D}_q^\nabla$. But then $d_q \xrightarrow{b} ts(d'_q)$ in $\mathcal{D}_q$ by definition.

The case of communication of $q$ with $p \neq r$ is similar to the above.

The last case is a communication between $q$ and $r$. So take $(d_q^1, (d_q^2, x))$ and suppose $(\pi(d_q^1), \pi(d_q^2, x)) \xrightarrow{b} (s'_q, s'_r)$ in $\mathcal{A}$. We have that $\pi^\nabla(d_q^1)$ is of the form $(s_q^1, a_1, s_r, f)$ and $\pi^\nabla(d_q^2)$ is of the form $(s_q^2, a_2, s_r, f)$; the $s_r$- and $f$-components are the same by Property (*). Moreover, by definition $\pi(d_q^1) = s_q^1$ holds. Let $s_r^1 = \pi(d_q^2, x)$, thus $s_r \xrightarrow{x} s_r^1$. These observations allow us to obtain the following sequence of transitions in $\mathcal{A}^\nabla$:

$$(s_q^1, a_1, s_r, f) \xrightarrow{x} (s_q^1, a_1, s_r^1, f|_x) \xrightarrow{b} (s'_q, s'_r)$$

Since $\mathcal{D}^\nabla$ satisfies the controllability condition we must have transitions $d_q^1 \xrightarrow{x} d'_q \xrightarrow{b} d''_q$ in $\mathcal{D}^\nabla$, with $\pi^\nabla(d''_q) = (s'_q, s'_r)$. This means that we have transition $(d_q^1, (d_q^2, x)) \xrightarrow{b} (ts(d''_q), (ts(d''_q), \varepsilon))$ in $\mathcal{D}$ and $\pi(ts(d''_q)) = s'_q$, $\pi(ts(d''_q), \varepsilon) = s'_r$. □

As $\mathcal{D}$ is covering, to prove that $\mathcal{D}$ is correct we need to show that all its maximal runs satisfy the correctness condition. For this we will construct for every run of $\mathcal{D}$ a corresponding run of $\mathcal{D}^\nabla$. The following definition and lemma tells us that it is enough to look at the runs of $\mathcal{D}$ of a special form.

**Definition 28 (*slow*)** We define $slow_r(\mathcal{D})$ as the set of all sequences labeling runs of $\mathcal{D}$ of the form $y_0x_0a_1 \cdots a_k y_k x_k a_{k+1} \ldots$ or $y_0x_0a_1 \cdots y_{k-1}x_{k-1}a_k x_k y_\omega$, where $a_i \in \Sigma_q \cap \Sigma_r$, $x_i \in (\Sigma_r^{loc})^*$, $y_i \in (\Sigma \setminus \Sigma_r)^*$, and $y_\omega \in (\Sigma \setminus \Sigma_r)^\omega$

**Lemma 29** A covering controller $\mathcal{D}$ is correct for $\mathcal{A}$ iff for all $w \in slow_r(\mathcal{D})$, $run(w)$ satisfies the correctness condition inherited from $\mathcal{A}$.

**Proof**
Observe first $\mathcal{D}$ is $r$-short, since $\mathcal{A}$ is $r$-short and $\mathcal{D}$ is covering. Thus every sequence labeling some run of $\mathcal{D}$ either has finitely many $r$-actions or infinitely many communications of $r$ with $q$.

Secondly, note that every sequence $w$ labeling some run of $\mathcal{D}$ can be rewritten into a sequence $w'$ from $slow_r(\mathcal{D})$ by repeatedly replacing factors $ab$ by $ba$, if $dom(a) \cap dom(b) = \emptyset$. We have that $run(w')$ is also defined and $run_p(w) =$

$run_p(w')$ for every process $p$. Therefore for correctness it will be enough to reason on sequences from $slow_r(\mathcal{D})$. □

For every sequence $w \in slow_r(\mathcal{D})$ as in Definition 28 we define the sequence $\chi(w) \in (\Sigma^\triangledown)^\infty$ by induction on the length of $w$. Let $\chi(\varepsilon) = ch(f_0)\, ch(a_0)$, where $f_0$ and $a_0$ are determined by the initial $q$-state of $\mathcal{D}^\triangledown$. For $w \in \Sigma^*, b \in \Sigma$ let

$$\chi(wb) = \begin{cases} \chi(w)b & \text{if } b \notin \Sigma_q \\ \chi(w)b\, ch(a) & \text{if } b \in \Sigma_q \setminus \Sigma_r \\ \chi(w)b\, ch(f)\, ch(a) & \text{if } b \in \Sigma_q \cap \Sigma_r. \end{cases}$$

where $a$ and $f$ are determined by the state reached by $\mathcal{D}^\triangledown$ on $\chi(w)b$. The next lemma implies the correctness of the construction, and at the same time confirms that the above definition makes sense, that is, the needed runs of $\mathcal{D}^\triangledown$ are defined.

**Lemma 30** For every sequence $w \in slow_r(\mathcal{D})$ we have that $run^\triangledown(\chi(w))$ is defined. If $w$ is finite then the states reached by $\mathcal{D}$ on $w$ and by $\mathcal{D}^\triangledown$ on $\chi(w)$ satisfy the following:

1. $state_p(w) = state_p^\triangledown(\chi(w))$ for every $p \neq q, r$.
2. Let $w = y_0x_0a_1 \cdots a_k y_k x_k$, where $a_i \in \Sigma_q \cap \Sigma_r$, $x_i \in (\Sigma_r^{loc})^*$, and $y_i \in (\Sigma \setminus \Sigma_r)^*$. Then $state_r(w) = (d_q, x_k)$ and $state_q(w) = d'_q$, where $d_q = state_q^\triangledown(\chi(y_0x_0a_1 \cdots a_k))$ and $d'_q = state_q^\triangledown(\chi(y_0x_0a_1 \cdots a_k y_k))$.

**Proof**
Induction on the length of $w = y_0x_0a_1 \cdots a_k y_k x_k$. If $w = \varepsilon$ then $state_q(\varepsilon) = d_q^1$ and $state_r(\varepsilon) = (d_q^1, \varepsilon)$ where $d_q^0 \xrightarrow{ch(f_0)\, ch(a)} d_q^1$ in $\mathcal{D}_q$, which shows the claim. Let $w = w'b$. If $b \notin (\Sigma_q \cup \Sigma_r)$, then $x_k = \varepsilon$, $y_k = y'b$, $\chi(w'b) = \chi(w')b$, $state_q(w'b) = state_q(w') \overset{ind.}{=} state_q^\triangledown(\chi(y_0x_0a_1 \cdots a_k y')) = state_q^\triangledown(\chi(y_0x_0a_1 \cdots a_k y')b)$. Moreover, $state_r(w'b) = state_r(w') \overset{ind.}{=} (d_q, \varepsilon)$, where $d_q = state_q^\triangledown(\chi(y_0x_0a_1 \cdots a_k))$. Finally, assuming that $run^\triangledown(\chi(w'))$ defined, observe that this run can be extended by a $b$-transition since it can be in $w$ and the concerned states are the same.

We consider the remaining cases:

1. Let $b \in \Sigma_r^{loc}$, then $\chi(w'b) = \chi(w')b$ and $x_k = x'b$. We have $state_q(w'b) = state_q(w') \overset{ind.}{=} state_q^\triangledown(\chi(y_0x_0a_1 \cdots y_k)) =: d'_q$. Moreover, $state_r(w') = (d_q, x')$, where $d_q = state_q^\triangledown(\chi(y_0x_0a_1 \cdots a_k))$. In $\mathcal{D}_r$ there is a transition $(d_q, x') \xrightarrow{b} (d_q, x'b)$, which shows the claim about states. Finally we justify that the run on $\chi(w')$ in $\mathcal{D}^\triangledown$ can be extended by a $b$. We know that $d_q \xrightarrow{x'b}$ and $d'_q \xrightarrow{x'}$ in $\mathcal{D}^\triangledown$, and want to show that $d'_q \xrightarrow{x'b}$. This holds since $\mathcal{D}^\triangledown$ is covering and since Property (*) guarantees that the $s_r$ and $f$ components of $\pi^\triangledown(d_q)$ and $\pi^\triangledown(d'_q)$ are the same.
2. Let $b \in \Sigma_q \setminus \Sigma_r$, so $b$ is either local on $q$ or a communication with $p \neq q, r$. We have $x_k = \varepsilon$ and $y_k = y'b$. Assume that $b$ is local on $q$. We have $\chi(w) = \chi(w')b\, ch(a)$, where $a \in \Sigma_q^{loc}$ and $d_q$ are such that $d_q = state_q^\triangledown(\chi(w'))$ and $d_q \xrightarrow{b} d_q^1 \xrightarrow{ch(a)} d_q^2$ in $\mathcal{D}^\triangledown$. By induction, $state_q(w') =$

$state_q^\nabla(\chi(y_0x_0a_1\cdots a_ky')) = d_q$, and by definition of $\mathcal{D}_q$, $d_q \xrightarrow{b} d_q^2 = ts(d_q^1)$. Thus $state_q(w) = d_q^2 = state_q^\nabla(\chi(w))$ and the claim about states is shown. The run on $\chi(w)$ in $\mathcal{D}^\nabla$ exists by the definition of $\chi(w)$ from $\chi(w')$.

The case of a communication with $p \neq r$ is similar to the above.

3. Let $b \in \Sigma_q \cap \Sigma_r$ be a communication between $q$ and $r$, thus $a_k = b$ and $x_k = y_k = \varepsilon$. We have $\chi(w) = \chi(w')b\,ch(f)\,ch(a)$, where $a, f$ are such that $state_q^\nabla(\chi(w')) = d_q \xrightarrow{b} d_q^1 \xrightarrow{ch(f)\,ch(a)} d_q^2$. Consider $d'_q = state_q^\nabla(\chi(y_0x_0a_1\cdots a_{k-1}))$ and $d''_q = state_q^\nabla(\chi(y_0x_0a_1\cdots a_{k-1}y_{k-1}))$. By induction, $state_q(w') = d''_q$ and $state_r(w') = (d'_q, x_{k-1})$. In $\mathcal{D}$ we have a transition $(d''_q, (d'_q, x_{k-1})) \xrightarrow{b} (d_q^2, (d_q^2, \varepsilon))$ since $d''_q \xrightarrow{x_{k-1}} d_q \xrightarrow{b} d_q^1 \xrightarrow{ch(f)\,ch(a)} d_q^2$ in $\mathcal{D}^\nabla$. Thus, $state_q(w) = d_q^2 = state_q^\nabla(\chi(w))$ and $state_r(w) = (d_q^2, \varepsilon) = (state_q^\nabla(\chi(w)), \varepsilon)$, which shows the claim about states. The run on $\chi(w)$ in $\mathcal{D}^\nabla$ exists by the definition of $\chi(w)$ from $\chi(w')$.

□

**Lemma 31** If $w \in slow_r(\mathcal{D})$ and $run(w)$ is maximal in $\mathcal{D}$, then $run(\chi(w))$ is maximal in $\mathcal{D}^\nabla$.

**Proof**
Recall first that $run(\chi(w))$ is not maximal only if for some finite prefix $x$ of $\chi(w)$, $run(x)$ can be extended by some action $a$ (and the processes in $dom(a)$ do not appear anymore in the remaining suffix of $\chi(w)$). From the definition of $\chi(w)$ it follows that it suffices to consider prefixes of $\chi(w)$ of the form $\chi(u)$, where $w = uv$ with $u$ finite. By Lemma 30 we note first that such an $a$ cannot be on processes other than $q$ or $r$, since $state_p(u) = state_p^\nabla(\chi(u))$ for all $p \neq q, r$.

We consider the remaining cases, and assume $u = y_0x_0a_1\cdots a_ky_kx_k$:

1. Assume that $\chi(u)$ can be extended by some $b \in \Sigma_r^{loc}$ in $\mathcal{D}^\nabla$, and let $d_q = state_q^\nabla(\chi(y_0x_0a_1\cdots a_k))$, $d'_q = state_q^\nabla(\chi(y_0x_0a_1\cdots a_ky_k))$, so $d'_q \xrightarrow{x_kb}$ in $\mathcal{D}_q^\nabla$. By Lemma 30 we have $state_r(u) = (d_q, x_k)$ and by Property (*), the $s_r$- and $f$-components of $\pi^\nabla(d_q)$ and $\pi^\nabla(d'_q)$ are the same. Since $\mathcal{D}^\nabla$ is covering, this means that $d_q \xrightarrow{x_kb}$, hence there is a run on $ubv$ in $\mathcal{D}$ so $w$ was not maximal.

2. Assume that $\chi(u)$ can be extended by some $b \in \Sigma_q \setminus \Sigma_r$ and recall from Lemma 30 that $state_q(u) = d_q$, where $d_q = state_q^\nabla(\chi(y_0x_0a_1\cdots a_ky_k))$. Consider $u_1 = y_0x_0a_1\cdots a_ky_kb$ and assume that $b$ is $q$-local (the case of a communication with $p \neq r$ is similar). We have $d_q \xrightarrow{x_k} d'_q \xrightarrow{b}$ in $\mathcal{D}^\nabla$ from some $d'_q$, and we want to show that $d_q \xrightarrow{b} d''_q$ for some $d''_q$. But this holds since $\mathcal{D}^\nabla$ is covering and the $s_q$ components of $\pi^\nabla(d_q)$ and $\pi^\nabla(d'_q)$ are the same. So the run of $w$ in $\mathcal{D}_q$ was not maximal, since there is a run on $ubv$ in $\mathcal{D}$.

3. Assume that $\chi(u)$ can be extended by some $b \in \Sigma_q \cap \Sigma_r$. Recall from Lemma 30 that $state_q(u) = d'_q$ and $state_r(u) = (d_q, x_k)$, where $d_q = state_q^\nabla(\chi(y_0x_0a_1\cdots a_k))$ and $d'_q = state_q^\nabla(\chi(y_0x_0a_1\cdots a_ky_k))$. We have

that $state_q^\nabla(\chi(u)) = d_q^1$ where $d'_q \xrightarrow{x_k} d_q^1$, and $d_q^1 \xrightarrow{b} d_q^2$. According to the definition of $\mathcal{D}$, there is a transition $(d'_q, (d_q, x_k)) \xrightarrow{b} (ts(d_q^2), (ts(d_q^2), \varepsilon))$ in $\mathcal{D}$, so that the run on $w$ was not maximal.

Note that a run on $\chi(u)$ cannot be extended by actions of the form $ch(a)$ or $ch(f)$, since $\mathcal{D}^\nabla$ is covering. So the above four cases exhaust all the possibilities. □

**Lemma 32** If $\mathcal{D}^\nabla$ is a correct covering controller for $\mathcal{A}^\nabla$, then $\mathcal{D}$ is a correct covering controller for $\mathcal{A}$.

**Proof**
By Lemma 29 it is enough to show that for all $w \in slow_r(\mathcal{D})$, $run(w)$ satisfies $Corr$. By Lemmas 30 and 31 the run on $\chi(w)$ exists and is maximal. Since $\mathcal{D}^\nabla$ is correct this run satisfies $Corr^\nabla$.

Consider a maximal run in $\mathcal{D}$, labeled by some $w \in slow_r(\mathcal{D})$. It is of one of the forms

$$y_0x_0a_1 \cdots a_k y_k x_k a_{k+1} \ldots \quad \text{or} \quad y_0x_0a_1 \cdots a_k x_k y_\omega$$

where $a_i \in \Sigma_q \cap \Sigma_r$, $x_i \in (\Sigma_r^{loc})^*$, $y_i \in (\Sigma \setminus \Sigma_r)^*$, and $y_\omega \in (\Sigma \setminus \Sigma_r)^\omega$

By Lemma 30 $run_p(w)$ and $run_p^\nabla(\chi(w))$ are the same for $p \neq q, r$. Since for such $p$ also the correctness conditions of $\mathcal{A}$ and $\mathcal{A}^\nabla$ are the same, and since $run_p^\nabla(\chi(w))$ satisfies $Corr_p^\nabla$, so does $run_p(w)$.

Considering $run_q(w)$, Lemma 30 gives us

$$state_q(y_0x_0a_1 \cdots a_k y_k) = state_q(y_0x_0a_1 \cdots a_k y_k x_k) = state_q^\nabla(\chi(y_0x_0a_1 \cdots a_k y_k))$$

for every $k$. Moreover, the $\mathcal{A}_q$-component does not change when going from $\pi^\nabla(\chi(y_0x_0a_1 \cdots a_k y_k))$ to $\pi^\nabla(\chi(y_0x_0a_1 \cdots a_k y_k x_k))$. Thus, $\pi(run_q(w))$ is equal to the projection on $\mathcal{A}_q$ of $\pi^\nabla(run_q(\chi(w)))$, so $run_q(w)$ satisfies $Corr_q$.

It remains to consider $run_r(w)$. For this we can use Lemma 30 obtaining $state_r(y_0x_0a_1 \cdots a_k y_k x_k) = (d_q, x_k)$ with $d_q = state_q^\nabla(\chi(y_0x_0a_1 \cdots a_k))$, for every $k$. Recall that $\pi(d_q, x)$ was defined as the $\mathcal{A}_r$-component of $\pi^\nabla(d'_q)$, where $d_q \xrightarrow{x} d'_q$ in $\mathcal{D}^\nabla$. Assume first that $w$ is of the form $y_0x_0a_1 \cdots a_k y_k x_k a_{k+1} \ldots$. Observe that $\pi(run_r(w))$ is equal to the projection on $\mathcal{A}_r$ of $\pi^\nabla(run_q^\nabla(\chi(w)))$, thus $run_r(w)$ satisfies $Corr_r$ because $run_r^\nabla(\chi(w))$ satisfies $Corr_r$. Let now $w$ be of the form $y_0x_0a_1 \cdots a_k x_k y_\omega$. Since $run(w)$ is maximal we have that $state_r(w) \in T_r$, again because $run_r^\nabla(\chi(w))$ satisfies $Corr_r$. □

# 4 Conclusion

We have considered a model obtained by instantiating Zielonka automata into the supervisory control framework of Ramadge and Wonham [15]. The result is a distributed synthesis framework that is both expressive and decidable in interesting cases. To substantiate we have sketched how to encode threaded boolean programs with compare-and-swap instructions. Our main decidability result (Theorem 7) shows that the synthesis problem is decidable for hierarchical architectures and for all local omega-regular specifications. Recall that in the Pnueli and Rosner setting essentially only pipeline architectures are decidable,

with an additional restriction that only the first and the last process in the pipeline can handle environment inputs. In our case all the process can interact with the environment.

The synthesis procedure presented here is in $k$-EXPTIME for architectures of depth $k$, in particular it is EXPTIME for the case of a one server communicating with clients who do not communicate between each other. From [5] we know that these bounds are tight.

This paper essentially closes the case of tree architectures introduced in [5]. The long standing open question is the decidability of the synthesis problem for all architectures [3].

# References


[1] A. Arnold, A. Vincent, and I. Walukiewicz. Games for synthesis of controllers with partial observation. *Theoretical Computer Science*, 303(1):7–34, 2003.

[2] B. Finkbeiner and S. Schewe. Uniform distributed synthesis. In *Proc. LICS* 2005.

[3] P. Gastin, B. Lerman, and M. Zeitoun. Distributed games with causal memory are decidable for series-parallel systems. In *Proc. FSTTCS* 2004.

[4] P. Gastin and N. Sznajder. Fair synthesis for asynchronous distributed systems. *ACM Transactions on Computational Logic*, 14(2): 9, 2013.

[5] B. Genest, H. Gimbert, A. Muscholl, and I. Walukiewicz. Asynchronous games over tree architectures. In *Proc. ICALP* 2013.

[6] S. Graf, D. Peled, and S. Quinton. Achieving distributed control through model checking. *Formal Methods in System Design*, 40(2):263–281, 2012.

[7] J. Gutierrez and G. Winskel. Borel determinacy of concurrent games. In *Proc. CONCUR* 2013.

[8] O. Kupferman and M. Vardi. Synthesizing distributed systems. In *Proc. LICS* 2001.

[9] P. Madhusudan and P. Thiagarajan. Distributed control and synthesis for local specifications. In *Proc. ICALP* 2001.

[10] P. Madhusudan, P. S. Thiagarajan, and S. Yang. The MSO theory of connectedly communicating processes. In *Proc. FSTTCS* 2005.

[11] P.-A. Melliès. Asynchronous games 2: The true concurrency of innocence. *TCS*, 358(2-3):200–228, 2006.

[12] M. Mukund and M. A. Sohoni. Keeping Track of the Latest Gossip in a Distributed System. *Distributed Computing*, 10(3):137–148, 1997.

[13] A. Muscholl and S. Schewe. Unlimited decidability of distributed synthesis with limited missing knowledge. In *Proc. MFCS* 2013.

[14] A. Pnueli and R. Rosner. Distributed reactive systems are hard to synthesize. In *Proc. FOCS* 1990.

[15] P. J. G. Ramadge and W. M. Wonham. The control of discrete event systems. *Proc. of the IEEE*, 77(2):81–98, 1989.

[16] I. Walukiewicz. Pushdown processes: Games and model checking. *Inf. Comput.*, 164(2):234–263, 2001.

[17] W. Zielonka. Notes on finite asynchronous automata. *RAIRO–Theoretical Informatics and Applications*, 21:99–135, 1987.